

VISIT...

LANZAROTE
Caliente.COM

part I

Introduction

The Science of Macroeconomics

The whole of science is nothing more than the refinement of everyday thinking.

— Albert Einstein

1-1 What Macroeconomists Study

Why have some countries experienced rapid growth in incomes over the past century while others stay mired in poverty? Why do some countries have high rates of inflation while others maintain stable prices? Why do all countries experience recessions and depressions—recurrent periods of falling incomes and rising unemployment—and how can government policy reduce the frequency and severity of these episodes? Macroeconomics, the study of the economy as a whole, attempts to answer these and many related questions.

To appreciate the importance of macroeconomics, you need only read the newspaper or listen to the news. Every day you can see headlines such as INCOME GROWTH SLOWS, FED MOVES TO COMBAT INFLATION, or STOCKS FALL AMID RECESSION FEARS. Although these macroeconomic events may seem abstract, they touch all of our lives. Business executives forecasting the demand for their products must guess how fast consumers' incomes will grow. Senior citizens living on fixed incomes wonder how fast prices will rise. Recent college graduates looking for jobs hope that the economy will boom and that firms will be hiring.

Because the state of the economy affects everyone, macroeconomic issues play a central role in political debate. Voters are aware of how the economy is doing, and they know that government policy can affect the economy in powerful ways. As a result, the popularity of the incumbent president rises when the economy is doing well and falls when it is doing poorly.

Macroeconomic issues are also at the center of world politics. In recent years, Europe has moved toward a common currency, many Asian countries have experienced financial turmoil and capital flight, and the United States has financed large trade deficits by borrowing from abroad. When world leaders meet, these topics are often high on their agendas.

Although the job of making economic policy falls to world leaders, the job of explaining how the economy as a whole works falls to macroeconomists. Toward this end, macroeconomists collect data on incomes, prices, unemployment, and many other variables from different time periods and different countries. They then attempt to formulate general theories that help to explain these data. Like astronomers studying the evolution of stars or biologists studying the evolution of species, macroeconomists cannot conduct controlled experiments. Instead, they must make use of the data that history gives them. Macroeconomists observe that economies differ from one another and that they change over time. These observations provide both the motivation for developing macroeconomic theories and the data for testing them.

To be sure, macroeconomics is a young and imperfect science. The macroeconomist's ability to predict the future course of economic events is no better than the meteorologist's ability to predict next month's weather. But, as you will see, macroeconomists do know quite a lot about how the economy works. This knowledge is useful both for explaining economic events and for formulating economic policy.

Every era has its own economic problems. In the 1970s, Presidents Richard Nixon, Gerald Ford, and Jimmy Carter all wrestled in vain with a rising rate of inflation. In the 1980s, inflation subsided, but Presidents Ronald Reagan and George Bush presided over large federal budget deficits. In the 1990s, with President Bill Clinton in the Oval Office, the budget deficit shrank and even turned into a budget surplus, but federal taxes as a share of national income reached a historic high. So it was no surprise that when President George W. Bush moved into the White House in 2001, he put a tax cut high on his agenda. The basic principles of macroeconomics do not change from decade to decade, but the macroeconomist must apply these principles with flexibility and creativity to meet changing circumstances.


CASE STUDY

The Historical Performance of the U.S. Economy

Economists use many types of data to measure the performance of an economy. Three macroeconomic variables are especially important: real gross domestic product (GDP), the inflation rate, and the unemployment rate. **Real GDP** measures the total income of everyone in the economy (adjusted for the level of prices). The **inflation rate** measures how fast prices are rising. The **unemployment rate** measures the fraction of the labor force that is out of work. Macroeconomists study how these variables are determined, why they change over time, and how they interact with one another.

Figure 1-1 shows real GDP per person in the United States. Two aspects of this figure are noteworthy. First, real GDP grows over time. Real GDP per person is today about five times its level in 1900. This growth in average

figure 1-1



Real GDP per Person in the U.S. Economy


Real GDP measures the total income of everyone in the economy, and real GDP per person measures the income of the average person in the economy. This figure shows that real GDP per person tends to grow over time and that this normal growth is sometimes interrupted by periods of declining income, called recessions or depressions.

Note: Real GDP is plotted here on a logarithmic scale. On such a scale, equal distances on the vertical axis represent equal *percentage* changes. Thus, the distance between \$5,000 and \$10,000 (a 100 percent change) is the same as the distance between \$10,000 and \$20,000 (a 100 percent change).
Source: U.S. Bureau of the Census (*Historical Statistics of the United States: Colonial Times to 1970*) and U.S. Department of Commerce.

income allows us to enjoy a higher standard of living than our great-grandparents did. Second, although real GDP rises in most years, this growth is not steady. There are repeated periods during which real GDP falls, the most dramatic instance being the early 1930s. Such periods are called **recessions** if they are mild and **depressions** if they are more severe. Not surprisingly, periods of declining income are associated with substantial economic hardship.

Figure 1-2 shows the U.S. inflation rate. You can see that inflation varies substantially. In the first half of the twentieth century, the inflation rate averaged only slightly above zero. Periods of falling prices, called **deflation**, were almost as common as periods of rising prices. In the past half century, inflation has been the norm. The inflation problem became most severe during the late 1970s, when prices rose at a rate of almost 10 percent per year. In recent years,

figure 1-2



The Inflation Rate in the U.S. Economy

The inflation rate measures the percentage change in the average level of prices from the year before. When the inflation rate is above zero, prices are rising. When it is below zero, prices are falling. If the inflation rate declines but remains positive, prices are rising but at a slower rate.

Note: The inflation rate is measured here using the GDP deflator.


Source: U.S. Bureau of the Census (*Historical Statistics of the United States: Colonial Times to 1970*) and U.S. Department of Commerce.

the inflation rate has been about 2 or 3 percent per year, indicating that prices have been fairly stable.

Figure 1-3 shows the U.S. unemployment rate. Notice that there is always some unemployment in our economy. In addition, although there is no long-term trend, the amount of unemployment varies from year to year. Recessions and depressions are associated with unusually high unemployment. The highest rates of unemployment were reached during the Great Depression of the 1930s.

These three figures offer a glimpse at the history of the U.S. economy. In the chapters that follow, we first discuss how these variables are measured and then develop theories to explain how they behave.

figure 1-3



The Unemployment Rate in the U.S. Economy

The unemployment rate measures the percentage of people in the labor force who do not have jobs. This figure shows that the economy always has some unemployment and that the amount fluctuates from year to year.

Source: U.S. Bureau of the Census (*Historical Statistics of the United States: Colonial Times to 1970*) and U.S. Department of Commerce.

1-2 How Economists Think

Although economists often study politically charged issues, they try to address these issues with a scientist's objectivity. Like any science, economics has its own set of tools—terminology, data, and a way of thinking—that can seem foreign and arcane to the layman. The best way to become familiar with these tools is to practice using them, and this book will afford you ample opportunity to do so. To make these tools less forbidding, however, let's discuss a few of them here.

Theory as Model Building

Young children learn much about the world around them by playing with toy versions of real objects. For instance, they often put together models of cars, trains, or planes. These models are far from realistic, but the model-builder learns

a lot from them nonetheless. The model illustrates the essence of the real object it is designed to resemble.

Economists also use **models** to understand the world, but an economist's model is more likely to be made of symbols and equations than plastic and glue. Economists build their “toy economies” to help explain economic variables, such as GDP, inflation, and unemployment. Economic models illustrate, often in mathematical terms, the relationships among the variables. They are useful because they help us to dispense with irrelevant details and to focus on important connections.

Models have two kinds of variables: endogenous variables and exogenous variables. **Endogenous variables** are those variables that a model tries to explain. **Exogenous variables** are those variables that a model takes as given. The purpose of a model is to show how the exogenous variables affect the endogenous variables. In other words, as Figure 1-4 illustrates, exogenous variables come from outside the model and serve as the model's input, whereas endogenous variables are determined inside the model and are the model's output.


To make these ideas more concrete, let's review the most celebrated of all economic models—the model of supply and demand. Imagine that an economist were interested in figuring out what factors influence the price of pizza and the quantity of pizza sold. He or she would develop a model that described the behavior of pizza buyers, the behavior of pizza sellers, and their interaction in the market for pizza. For example, the economist supposes that the quantity of pizza demanded by consumers Q^d depends on the price of pizza P and on aggregate income Y . This relationship is expressed in the equation

$$Q^d = D(P, Y),$$

where $D()$ represents the demand function. Similarly, the economist supposes that the quantity of pizza supplied by pizzerias Q^s depends on the price of pizza P and on the price of materials P_m , such as cheese, tomatoes, flour, and anchovies. This relationship is expressed as

$$Q^s = S(P, P_m),$$

figure 1-4



How Models Work

Models are simplified theories that show the key relationships among economic variables. The exogenous variables are those that come from outside the model. The endogenous variables are those that the model explains. The model shows how changes in the exogenous variables affect the endogenous variables.

where $S()$ represents the supply function. Finally, the economist assumes that the price of pizza adjusts to bring the quantity supplied and quantity demanded into balance:


$$Q^s = Q^d.$$

These three equations compose a model of the market for pizza.

The economist illustrates the model with a supply-and-demand diagram, as in Figure 1-5. The demand curve shows the relationship between the quantity of pizza demanded and the price of pizza, while holding aggregate income constant. The demand curve slopes downward because a higher price of pizza encourages consumers to switch to other foods and buy less pizza. The supply curve shows the relationship between the quantity of pizza supplied and the price of pizza, while holding the price of materials constant. The supply curve slopes upward because a higher price of pizza makes selling pizza more profitable, which encourages pizzerias to produce more of it. The equilibrium for the market is the price and quantity at which the supply and demand curves intersect. At the equilibrium price, consumers choose to buy the amount of pizza that pizzerias choose to produce.

This model of the pizza market has two exogenous variables and two endogenous variables. The exogenous variables are aggregate income and the price of materials. The model does not attempt to explain them but takes them as given (perhaps to be explained by another model). The endogenous variables are the

figure 1-5




The Model of Supply and Demand

The most famous economic model is that of supply and demand for a good or service—in this case, pizza. The demand curve is a downward-sloping curve relating the price of pizza to the quantity of pizza that consumers demand. The supply curve is an upward-sloping curve relating the price of pizza to the quantity of pizza that pizzerias supply. The price of pizza adjusts until the quantity supplied equals the quantity demanded. The point where the two curves cross is the market equilibrium, which shows the equilibrium price of pizza and the equilibrium quantity of pizza.


price of pizza and the quantity of pizza exchanged. These are the variables that the model attempts to explain.

The model can be used to show how a change in one of the exogenous variables affects both endogenous variables. For example, if aggregate income increases, then the demand for pizza increases, as in panel (a) of Figure 1-6. The model shows that both the equilibrium price and the equilibrium quantity of pizza rise. Similarly, if the price of materials increases, then the supply of pizza decreases, as in panel (b) of Figure 1-6. The model shows that in this case the equilibrium price of pizza rises and the equilibrium quantity of pizza falls. Thus, the model shows how changes in aggregate income or in the price of materials affect price and quantity in the market for pizza.

figure 1-6



(b) A Shift in Supply



Changes in Equilibrium

In panel (a), a rise in aggregate income causes the demand for pizza to increase: at any given price, consumers now want to buy more pizza. This is represented by a rightward shift in the demand curve from D_1 to D_2 . The market moves to the new intersection of supply and demand. The equilibrium price rises from P_1 to P_2 , and the equilibrium quantity of pizza rises from Q_1 to Q_2 . In panel (b), a rise in the price of materials decreases the supply of pizza: at any given price, pizzerias find that the sale of pizza is less profitable and therefore choose to produce less pizza. This is represented by a leftward shift in the supply curve from S_1 to S_2 . The market moves to the new intersection of supply and demand. The equilibrium price rises from P_1 to P_2 , and the equilibrium quantity falls from Q_1 to Q_2 .

Like all models, this model of the pizza market makes simplifying assumptions. The model does not take into account, for example, that every pizzeria is in a different location. For each customer, one pizzeria is more convenient than the others, and thus pizzerias have some ability to set their own prices. Although the model assumes that there is a single price for pizza, in fact there could be a different price at every pizzeria.

How should we react to the model's lack of realism? Should we discard the simple model of pizza supply and pizza demand? Should we attempt to build a more complex model that allows for diverse pizza prices? The answers to these questions depend on our purpose. If our goal is to explain how the price of cheese affects the average price of pizza and the amount of pizza sold, then the diversity of pizza prices is probably not important. The simple model of the pizza market does a good job of addressing that issue. Yet if our goal is to explain why towns with three pizzerias have lower pizza prices than towns with one pizzeria, the simple model is less useful.

FYI

Using Functions to Express Relationships Among Variables

All economic models express relationships among economic variables. Often, these relationships are expressed as functions. A *function* is a mathematical concept that shows how one variable depends on a set of other variables. For example, in the model of the pizza market, we said that the quantity of pizza demanded depends on the price of pizza and on aggregate income. To express this, we use functional notation to write

$$Q^d = D(P, Y).$$

This equation says that the quantity of pizza demanded Q^d is a function of the price of pizza P and aggregate income Y . In functional notation, the variable preceding the parentheses denotes the function. In this case, $D()$ is the function expressing how the variables in parentheses determine the quantity of pizza demanded.

If we knew more about the pizza market, we could give a numerical formula for the quantity of pizza demanded. We might be able to write

$$Q^d = 60 - 10P + 2Y.$$

In this case, the demand function is

$$D(P, Y) = 60 - 10P + 2Y.$$

For any price of pizza and aggregate income, this function gives the corresponding quantity of pizza demanded. For example, if aggregate income is \$10 and the price of pizza is \$2, then the quantity of pizza demanded is 60 pies; if the price of pizza rises to \$3, the quantity of pizza demanded falls to 50 pies.

Functional notation allows us to express a relationship among variables even when the precise numerical relationship is unknown. For example, we might know that the quantity of pizza demanded falls when the price rises from \$2 to \$3, but we might not know by how much it falls. In this case, functional notation is useful: as long as we know that a relationship among the variables exists, we can remind ourselves of that relationship using functional notation.

The art in economics is in judging when an assumption is clarifying and when it is misleading. Any model constructed to be completely realistic would be too complicated for anyone to understand. Simplification is a necessary part of building a useful model. Yet models lead to incorrect conclusions if they assume away features of the economy that are crucial to the issue at hand. Economic modeling therefore requires care and common sense.

A Multitude of Models

Macroeconomists study many facets of the economy. For example, they examine the role of saving in economic growth, the impact of labor unions on unemployment, the effect of inflation on interest rates, and the influence of trade policy on the trade balance and exchange rates. Macroeconomics is as diverse as the economy.

Although economists use models to address all these issues, no single model can answer all questions. Just as carpenters use different tools for different tasks, economists use different models to explain different economic phenomena. Students of macroeconomics, therefore, must keep in mind that there is no single “correct” model useful for all purposes. Instead, there are many models, each of which is useful for shedding light on a different facet of the economy. The field of macroeconomics is like a Swiss army knife—a set of complementary but distinct tools that can be applied in different ways in different circumstances.

This book therefore presents many different models that address different questions and that make different assumptions. Remember that a model is only as good as its assumptions and that an assumption that is useful for some purposes may be misleading for others. When using a model to address a question, the economist must keep in mind the underlying assumptions and judge whether these are reasonable for the matter at hand.

Prices: Flexible Versus Sticky

Throughout this book, one group of assumptions will prove especially important—those concerning the speed with which wages and prices adjust. Economists normally presume that the price of a good or a service moves quickly to bring quantity supplied and quantity demanded into balance. In other words, they assume that a market goes to the equilibrium of supply and demand. This assumption is called **market clearing** and is central to the model of the pizza market discussed earlier. For answering most questions, economists use market-clearing models.

Yet the assumption of *continuous* market clearing is not entirely realistic. For markets to clear continuously, prices must adjust instantly to changes in supply and demand. In fact, however, many wages and prices adjust slowly. Labor contracts often set wages for up to three years. Many firms leave their product prices the same for long periods of time—for example, magazine publishers typically

change their newsstand prices only every three or four years. Although market-clearing models assume that all wages and prices are **flexible**, in the real world some wages and prices are **sticky**.

The apparent stickiness of prices does not make market-clearing models useless. After all, prices are not stuck forever; eventually, they do adjust to changes in supply and demand. Market-clearing models might not describe the economy at every instant, but they do describe the equilibrium toward which the economy gravitates. Therefore, most macroeconomists believe that price flexibility is a good assumption for studying long-run issues, such as the growth in real GDP that we observe from decade to decade.

For studying short-run issues, such as year-to-year fluctuations in real GDP and unemployment, the assumption of price flexibility is less plausible. Over short periods, many prices are fixed at predetermined levels. Therefore, most macroeconomists believe that price stickiness is a better assumption for studying the behavior of the economy in the short run.

Microeconomic Thinking and Macroeconomic Models

Microeconomics is the study of how households and firms make decisions and how these decisionmakers interact in the marketplace. A central principle of microeconomics is that households and firms *optimize*—they do the best they can for themselves given their objectives and the constraints they face. In microeconomic models, households choose their purchases to maximize their level of satisfaction, which economists call *utility*, and firms make production decisions to maximize their profits.

Because economy-wide events arise from the interaction of many households and many firms, macroeconomics and microeconomics are inextricably linked. When we study the economy as a whole, we must consider the decisions of individual economic actors. For example, to understand what determines total consumer spending, we must think about a family deciding how much to spend today and how much to save for the future. To understand what determines total investment spending, we must think about a firm deciding whether to build a new factory. Because aggregate variables are the sum of the variables describing many individual decisions, macroeconomic theory rests on a microeconomic foundation.

Although microeconomic decisions always underlie economic models, in many models the optimizing behavior of households and firms is implicit rather than explicit. The model of the pizza market we discussed earlier is an example. Households' decisions about how much pizza to buy underlie the demand for pizza, and pizzerias' decisions about how much pizza to produce underlie the supply of pizza. Presumably, households make their decisions to maximize utility, and pizzerias make their decisions to maximize profit. Yet the model did not focus on these microeconomic decisions; it left them in the background. Similarly, in much of macroeconomics, the optimizing behavior of households and firms is left implicit.

1-3 How This Book Proceeds

This book has six parts. This chapter and the next make up Part One, the Introduction. Chapter 2 discusses how economists measure economic variables, such as aggregate income, the inflation rate, and the unemployment rate.

Part Two, “Classical Theory: The Economy in the Long Run,” presents the classical model of how the economy works. The key assumption of the classical model is that prices are flexible. That is, with rare exceptions, the classical model assumes market clearing. Because the assumption of price flexibility describes the economy only in the long run, classical theory is best suited for analyzing a time horizon of at least several years.

Part Three, “Growth Theory: The Economy in the Very Long Run,” builds on the classical model. It maintains the assumption of market clearing but adds a new emphasis on growth in the capital stock, the labor force, and technological knowledge. Growth theory is designed to explain how the economy evolves over a period of several decades.

Part Four, “Business Cycle Theory: The Economy in the Short Run,” examines the behavior of the economy when prices are sticky. The non-market-clearing model developed here is designed to analyze short-run issues, such as the reasons for economic fluctuations and the influence of government policy on those fluctuations. It is best suited to analyzing the changes in the economy we observe from month to month or from year to year.

Part Five, “Macroeconomic Policy Debates,” builds on the previous analysis to consider what role the government should take in the economy. It considers how, if at all, the government should respond to short-run fluctuations in real GDP and unemployment. It also examines the various views on the effects of government debt.

Part Six, “More on the Microeconomics Behind Macroeconomics,” presents some of the microeconomic models that are useful for analyzing macroeconomic issues. For example, it examines the household’s decisions regarding how much to consume and how much money to hold and the firm’s decision regarding how much to invest. These individual decisions together form the larger macroeconomic picture. The goal of studying these microeconomic decisions in detail is to refine our understanding of the aggregate economy.

Summary

1. Macroeconomics is the study of the economy as a whole—including growth in incomes, changes in prices, and the rate of unemployment. Macroeconomists attempt both to explain economic events and to devise policies to improve economic performance.
2. To understand the economy, economists use models—theories that simplify reality in order to reveal how exogenous variables influence endogenous variables. The art in the science of economics is in judging whether a model

captures the important economic relationships for the matter at hand. Because no single model can answer all questions, macroeconomists use different models to look at different issues.

3. A key feature of a macroeconomic model is whether it assumes that prices are flexible or sticky. According to most macroeconomists, models with flexible prices describe the economy in the long run, whereas models with sticky prices offer a better description of the economy in the short run.
4. Microeconomics is the study of how firms and individuals make decisions and how these decisionmakers interact. Because macroeconomic events arise from many microeconomic interactions, macroeconomists use many of the tools of microeconomics.

KEY CONCEPTS

Macroeconomics	Depression	Exogenous variables
Real GDP	Deflation	Market clearing
Inflation rate	Models	Flexible and sticky prices
Unemployment rate	Endogenous variables	Microeconomics
Recession		

QUESTIONS FOR REVIEW

1. Explain the difference between macroeconomics and microeconomics. How are these two fields related?
2. Why do economists build models?
3. What is a market-clearing model? When is the assumption of market clearing appropriate?

PROBLEMS AND APPLICATIONS

1. What macroeconomic issues have been in the news lately?
2. What do you think are the defining characteristics of a science? Does the study of the economy have these characteristics? Do you think macroeconomics should be called a science? Why or why not?
3. Use the model of supply and demand to explain how a fall in the price of frozen yogurt would affect the price of ice cream and the quantity of ice cream sold. In your explanation, identify the exogenous and endogenous variables.
4. How often does the price you pay for a haircut change? What does your answer imply about the usefulness of market-clearing models for analyzing the market for haircuts?

The Data of Macroeconomics

It is a capital mistake to theorize before one has data. Insensibly one begins to twist facts to suit theories, instead of theories to fit facts.

— Sherlock Holmes

Scientists, economists, and detectives have much in common: they all want to figure out what's going on in the world around them. To do this, they rely on both theory and observation. They build theories in an attempt to make sense of what they see happening. They then turn to more systematic observation to evaluate the theories' validity. Only when theory and evidence come into line do they feel they understand the situation. This chapter discusses the types of observation that economists use to develop and test their theories.

Casual observation is one source of information about what's happening in the economy. When you go shopping, you see how fast prices are rising. When you look for a job, you learn whether firms are hiring. Because we are all participants in the economy, we get some sense of economic conditions as we go about our lives.

A century ago, economists monitoring the economy had little more to go on than these casual observations. Such fragmentary information made economic policymaking all the more difficult. One person's anecdote would suggest the economy was moving in one direction, while a different person's anecdote would suggest it was moving in another. Economists needed some way to combine many individual experiences into a coherent whole. There was an obvious solution: as the old quip goes, the plural of "anecdote" is "data."

Today, economic data offer a systematic and objective source of information, and almost every day the newspaper has a story about some newly released statistic. Most of these statistics are produced by the government. Various government agencies survey households and firms to learn about their economic activity—how much they are earning, what they are buying, what prices they are charging, whether they have a job or are looking for work, and so on. From these surveys, various statistics are computed that summarize the state of the economy. Economists use these statistics to study the economy; policymakers use them to monitor developments and formulate policies.

This chapter focuses on the three statistics that economists and policymakers use most often. **Gross domestic product**, or **GDP**, tells us the nation's total income

and the total expenditure on its output of goods and services. The **consumer price index**, or **CPI**, measures the level of prices. The **unemployment rate** tells us the fraction of workers who are unemployed. In the following pages, we see how these statistics are computed and what they tell us about the economy.

2-1 Measuring the Value of Economic Activity: Gross Domestic Product

Gross domestic product is often considered the best measure of how well the economy is performing. This statistic is computed every three months by the Bureau of Economic Analysis (a part of the U.S. Department of Commerce) from a large number of primary data sources. The goal of GDP is to summarize in a single number the dollar value of economic activity in a given period of time.

There are two ways to view this statistic. One way to view GDP is as *the total income of everyone in the economy*. Another way to view GDP is as *the total expenditure on the economy's output of goods and services*. From either viewpoint, it is clear why GDP is a gauge of economic performance. GDP measures something people care about—their incomes. Similarly, an economy with a large output of goods and services can better satisfy the demands of households, firms, and the government.

How can GDP measure both the economy's income and the expenditure on its output? The reason is that these two quantities are really the same: for the economy as a whole, income must equal expenditure. That fact, in turn, follows from an even more fundamental one: because every transaction has both a buyer and a seller, every dollar of expenditure by a buyer must become a dollar of income to a seller. When Joe paints Jane's house for \$1,000, that \$1,000 is income to Joe and expenditure by Jane. The transaction contributes \$1,000 to GDP, regardless of whether we are adding up all income or adding up all expenditure.

To understand the meaning of GDP more fully, we turn to **national income accounting**, the accounting system used to measure GDP and many related statistics.


Income, Expenditure, and the Circular Flow

Imagine an economy that produces a single good, bread, from a single input, labor. Figure 2-1 illustrates all the economic transactions that occur between households and firms in this economy.

The inner loop in Figure 2-1 represents the flows of bread and labor. The households sell their labor to the firms. The firms use the labor of their workers to produce bread, which the firms in turn sell to the households. Hence, labor flows from households to firms, and bread flows from firms to households.

The outer loop in Figure 2-1 represents the corresponding flow of dollars. The households buy bread from the firms. The firms use some of the revenue

figure 2-1



The Circular Flow This figure illustrates the flows between firms and households in an economy that produces one good, bread, from one input, labor. The inner loop represents the flows of labor and bread: households sell their labor to firms, and the firms sell the bread they produce to households. The outer loop represents the corresponding flows of dollars: households pay the firms for the bread, and the firms pay wages and profit to the households. In this economy, GDP is both the total expenditure on bread and the total income from the production of bread.

from these sales to pay the wages of their workers, and the remainder is the profit belonging to the owners of the firms (who themselves are part of the household sector). Hence, expenditure on bread flows from households to firms, and income in the form of wages and profit flows from firms to households.

GDP measures the flow of dollars in this economy. We can compute it in two ways. GDP is the total income from the production of bread, which equals the sum of wages and profit—the top half of the circular flow of dollars. GDP is also the total expenditure on purchases of bread—the bottom half of the circular flow of dollars. To compute GDP, we can look at either the flow of dollars from firms to households or the flow of dollars from households to firms.

These two ways of computing GDP must be equal because the expenditure of buyers on products is, by the rules of accounting, income to the sellers of those products. Every transaction that affects expenditure must affect income, and every transaction that affects income must affect expenditure. For example, suppose that a firm produces and sells one more loaf of bread to a household. Clearly this transaction raises total expenditure on bread, but it also has an equal effect on total income. If the firm produces the extra loaf without hiring any more labor (such as by making the production process more efficient), then profit increases. If the firm produces the extra loaf by hiring more labor, then wages increase. In both cases, expenditure and income increase equally.

FYI

Stocks and Flows

Many economic variables measure a quantity of something—a quantity of money, a quantity of goods, and so on. Economists distinguish between two types of quantity variables: stocks and flows. A **stock** is a quantity measured at a given point in time, whereas a **flow** is a quantity measured per unit of time.

The bathtub, shown in Figure 2-2, is the classic example used to illustrate stocks and flows. The amount of water in the tub is a stock: it is the quantity of water in the tub at a given point in time. The amount of water coming out of the faucet is a flow: it is the quantity of water being added to the tub per unit of time. Note that we measure stocks and flows in different units. We say that the bathtub contains 50 *gallons* of water, but that water is coming out of the faucet at 5 *gallons per minute*.

GDP is probably the most important flow variable in economics: it tells us how many dollars are flowing around the economy's circular flow per unit of time. When you hear someone say that the U.S. GDP is \$10 trillion, you should


understand that this means that it is \$10 trillion *per year*. (Equivalently, we could say that U.S. GDP is \$317,000 per second.)

Stocks and flows are often related. In the bathtub example, these relationships are clear. The stock of water in the tub represents the accumulation of the flow out of the faucet, and the flow of water represents the change in the stock. When building theories to explain economic variables, it is often useful to determine whether the variables are stocks or flows and whether any relationships link them.

Here are some examples of related stocks and flows that we study in future chapters:

- A person's wealth is a stock; income and expenditure are flows.
- The number of unemployed people is a stock; the number of people losing their jobs is a flow.
- The amount of capital in the economy is a stock; the amount of investment is a flow.
- The government debt is a stock; the government budget deficit is a flow.

figure 2-2



Stocks and Flows The amount of water in a bathtub is a stock: it is a quantity measured at a given moment in time. The amount of water coming out of the faucet is a flow: it is a quantity measured per unit of time.

Rules for Computing GDP

In an economy that produces only bread, we can compute GDP by adding up the total expenditure on bread. Real economies, however, include the production and sale of a vast number of goods and services. To compute GDP for such a complex economy, it will be helpful to have a more precise definition: *gross*

domestic product (GDP) is the market value of all final goods and services produced within an economy in a given period of time. To see how this definition is applied, let's discuss some of the rules that economists follow in constructing this statistic.

Adding Apples and Oranges The U.S. economy produces many different goods and services—hamburgers, haircuts, cars, computers, and so on. GDP combines the value of these goods and services into a single measure. The diversity of products in the economy complicates the calculation of GDP because different products have different values.

Suppose, for example, that the economy produces four apples and three oranges. How do we compute GDP? We could simply add apples and oranges and conclude that GDP equals seven pieces of fruit. But this makes sense only if we thought apples and oranges had equal value, which is generally not true. (This would be even clearer if the economy had produced four watermelons and three grapes.)

To compute the total value of different goods and services, the national income accounts use market prices because these prices reflect how much people are willing to pay for a good or service. Thus, if apples cost \$0.50 each and oranges cost \$1.00 each, GDP would be

$$\begin{aligned}\text{GDP} &= (\text{Price of Apples} \times \text{Quantity of Apples}) \\ &\quad + (\text{Price of Oranges} \times \text{Quantity of Oranges}) \\ &= (\$0.50 \times 4) + (\$1.00 \times 3) \\ &= \$5.00.\end{aligned}$$

GDP equals \$5.00—the value of all the apples, \$2.00, plus the value of all the oranges, \$3.00.

Used Goods When the Topps Company makes a package of baseball cards and sells it for 50 cents, that 50 cents is added to the nation's GDP. But what about when a collector sells a rare Mickey Mantle card to another collector for \$500? That \$500 is not part of GDP. GDP measures the value of currently produced goods and services. The sale of the Mickey Mantle card reflects the transfer of an asset, not an addition to the economy's income. Thus, the sale of used goods is not included as part of GDP.

The Treatment of Inventories Imagine that a bakery hires workers to produce more bread, pays their wages, and then fails to sell the additional bread. How does this transaction affect GDP?

The answer depends on what happens to the unsold bread. Let's first suppose that the bread spoils. In this case, the firm has paid more in wages but has not received any additional revenue, so the firm's profit is reduced by the amount that wages are increased. Total expenditure in the economy hasn't changed because no one buys the bread. Total income hasn't changed either—although more is distributed as wages and less as profit. Because the transaction affects neither expenditure nor income, it does not alter GDP.

Now suppose, instead, that the bread is put into inventory to be sold later. In this case, the transaction is treated differently. The owners of the firm are assumed to have “purchased” the bread for the firm's inventory, and the firm's profit is not

reduced by the additional wages it has paid. Because the higher wages raise total income, and greater spending on inventory raises total expenditure, the economy's GDP rises.

What happens later when the firm sells the bread out of inventory? This case is much like the sale of a used good. There is spending by bread consumers, but there is inventory disinvestment by the firm. This negative spending by the firm offsets the positive spending by consumers, so the sale out of inventory does not affect GDP.

The general rule is that when a firm increases its inventory of goods, this investment in inventory is counted as expenditure by the firm owners. Thus, production for inventory increases GDP just as much as production for final sale. A sale out of inventory, however, is a combination of positive spending (the purchase) and negative spending (inventory disinvestment), so it does not influence GDP. This treatment of inventories ensures that GDP reflects the economy's current production of goods and services.

Intermediate Goods and Value Added Many goods are produced in stages: raw materials are processed into intermediate goods by one firm and then sold to another firm for final processing. How should we treat such products when computing GDP? For example, suppose a cattle rancher sells one-quarter pound of meat to McDonald's for \$0.50, and then McDonald's sells you a hamburger for \$1.50. Should GDP include both the meat and the hamburger (a total of \$2.00), or just the hamburger (\$1.50)?

The answer is that GDP includes only the value of final goods. Thus, the hamburger is included in GDP but the meat is not: GDP increases by \$1.50, not by \$2.00. The reason is that the value of intermediate goods is already included as part of the market price of the final goods in which they are used. To add the intermediate goods to the final goods would be double counting—that is, the meat would be counted twice. Hence, GDP is the total value of final goods and services produced.

One way to compute the value of all final goods and services is to sum the value added at each stage of production. The **value added** of a firm equals the value of the firm's output less the value of the intermediate goods that the firm purchases. In the case of the hamburger, the value added of the rancher is \$0.50 (assuming that the rancher bought no intermediate goods), and the value added of McDonald's is $\$1.50 - \0.50 , or \$1.00. Total value added is $\$0.50 + \1.00 , which equals \$1.50. For the economy as a whole, the sum of all value added must equal the value of all final goods and services. Hence, GDP is also the total value added of all firms in the economy.

Housing Services and Other Imputations Although most goods and services are valued at their market prices when computing GDP, some are not sold in the marketplace and therefore do not have market prices. If GDP is to include the value of these goods and services, we must use an estimate of their value. Such an estimate is called an **imputed value**.

Imputations are especially important for determining the value of housing. A person who rents a house is buying housing services and providing income for the landlord; the rent is part of GDP, both as expenditure by the renter and as income for the landlord. Many people, however, live in their own homes. Although they do not pay rent to a landlord, they are enjoying housing services similar to those that

renters purchase. To take account of the housing services enjoyed by homeowners, GDP includes the “rent” that these homeowners “pay” to themselves. Of course, homeowners do not in fact pay themselves this rent. The Department of Commerce estimates what the market rent for a house would be if it were rented and includes that imputed rent as part of GDP. This imputed rent is included both in the homeowner’s expenditure and in the homeowner’s income.

Imputations also arise in valuing government services. For example, police officers, firefighters, and senators provide services to the public. Giving a value to these services is difficult because they are not sold in a marketplace and therefore do not have a market price. The national income accounts include these services in GDP by valuing them at their cost. That is, the wages of these public servants are used as a measure of the value of their output.

In many cases, an imputation is called for in principle but, to keep things simple, is not made in practice. Because GDP includes the imputed rent on owner-occupied houses, one might expect it also to include the imputed rent on cars, lawn mowers, jewelry, and other durable goods owned by households. Yet the value of these rental services is left out of GDP. In addition, some of the output of the economy is produced and consumed at home and never enters the marketplace. For example, meals cooked at home are similar to meals cooked at a restaurant, yet the value added in meals at home is left out of GDP.

Finally, no imputation is made for the value of goods and services sold in the *underground economy*. The underground economy is the part of the economy that people hide from the government either because they wish to evade taxation or because the activity is illegal. Domestic workers paid “off the books” is one example. The illegal drug trade is another.

Because the imputations necessary for computing GDP are only approximate, and because the value of many goods and services is left out altogether, GDP is an imperfect measure of economic activity. These imperfections are most problematic when comparing standards of living across countries. The size of the underground economy, for instance, varies from country to country. Yet as long as the magnitude of these imperfections remains fairly constant over time, GDP is useful for comparing economic activity from year to year.

Real GDP Versus Nominal GDP

Economists use the rules just described to compute GDP, which values the economy’s total output of goods and services. But is GDP a good measure of economic well-being? Consider once again the economy that produces only apples and oranges. In this economy GDP is the sum of the value of all the apples produced and the value of all the oranges produced. That is,

$$\begin{aligned} \text{GDP} = & (\text{Price of Apples} \times \text{Quantity of Apples}) \\ & + (\text{Price of Oranges} \times \text{Quantity of Oranges}). \end{aligned}$$

Notice that GDP can increase either because prices rise or because quantities rise.

It is easy to see that GDP computed this way is not a good gauge of economic well-being. That is, this measure does not accurately reflect how well the

economy can satisfy the demands of households, firms, and the government. If all prices doubled without any change in quantities, GDP would double. Yet it would be misleading to say that the economy's ability to satisfy demands has doubled, because the quantity of every good produced remains the same. Economists call the value of goods and services measured at current prices **nominal GDP**.

A better measure of economic well-being would tally the economy's output of goods and services and would not be influenced by changes in prices. For this purpose, economists use **real GDP**, which is the value of goods and services measured using a constant set of prices. That is, real GDP shows what would have happened to expenditure on output if quantities had changed but prices had not.

To see how real GDP is computed, imagine we wanted to compare output in 2002 and output in 2003 in our apple-and-orange economy. We could begin by choosing a set of prices, called *base-year prices*, such as the prices that prevailed in 2002. Goods and services are then added up using these base-year prices to value the different goods in both years. Real GDP for 2002 would be

$$\begin{aligned} \text{Real GDP} = & (2002 \text{ Price of Apples} \times 2002 \text{ Quantity of Apples}) \\ & + (2002 \text{ Price of Oranges} \times 2002 \text{ Quantity of Oranges}). \end{aligned}$$

Similarly, real GDP in 2003 would be

$$\begin{aligned} \text{Real GDP} = & (2002 \text{ Price of Apples} \times 2003 \text{ Quantity of Apples}) \\ & + (2002 \text{ Price of Oranges} \times 2003 \text{ Quantity of Oranges}). \end{aligned}$$

And real GDP in 2004 would be

$$\begin{aligned} \text{Real GDP} = & (2002 \text{ Price of Apples} \times 2004 \text{ Quantity of Apples}) \\ & + (2002 \text{ Price of Oranges} \times 2004 \text{ Quantity of Oranges}). \end{aligned}$$

Notice that 2002 prices are used to compute real GDP for all three years. Because the prices are held constant, real GDP varies from year to year only if the quantities produced vary. Because a society's ability to provide economic satisfaction for its members ultimately depends on the quantities of goods and services produced, real GDP provides a better measure of economic well-being than nominal GDP.

The GDP Deflator

From nominal GDP and real GDP we can compute a third statistic: the GDP deflator. The **GDP deflator**, also called the implicit price deflator for GDP, is defined as the ratio of nominal GDP to real GDP:

$$\text{GDP Deflator} = \frac{\text{Nominal GDP}}{\text{Real GDP}}.$$

The GDP deflator reflects what's happening to the overall level of prices in the economy.

To better understand this, consider again an economy with only one good, bread. If P is the price of bread and Q is the quantity sold, then nominal GDP is

the total number of dollars spent on bread in that year, $P \times Q$. Real GDP is the number of loaves of bread produced in that year times the price of bread in some base year, $P_{\text{base}} \times Q$. The GDP deflator is the price of bread in that year relative to the price of bread in the base year, P/P_{base} .

The definition of the GDP deflator allows us to separate nominal GDP into two parts: one part measures quantities (real GDP) and the other measures prices (the GDP deflator). That is,

$$\text{Nominal GDP} = \text{Real GDP} \times \text{GDP Deflator}.$$

Nominal GDP measures the current dollar value of the output of the economy. Real GDP measures output valued at constant prices. The GDP deflator measures the price of output relative to its price in the base year. We can also write this equation as

$$\text{Real GDP} = \frac{\text{Nominal GDP}}{\text{GDP Deflator}}.$$

In this form, you can see how the deflator earns its name: it is used to deflate (that is, take inflation out of) nominal GDP to yield real GDP.

Chain-Weighted Measures of Real GDP

We have been discussing real GDP as if the prices used to compute this measure never change from their base-year values. If this were truly the case, over time the prices would become more and more dated. For instance, the price of computers has fallen substantially in recent years, while the price of a year at college has risen. When valuing the production of computers and education, it would be misleading to use the prices that prevailed ten or twenty years ago.

To solve this problem, the Bureau of Economic Analysis used to update periodically the prices used to compute real GDP. About every five years, a new base year was chosen. The prices were then held fixed and used to measure year-to-year changes in the production of goods and services until the base year was updated once again.

In 1995, the bureau announced a new policy for dealing with changes in the base year. In particular, it now emphasizes *chain-weighted* measures of real GDP. With these new measures, the base year changes continuously over time. In essence, average prices in 2001 and 2002 are used to measure real growth from 2001 to 2002; average prices in 2002 and 2003 are used to measure real growth from 2002 to 2003; and so on. These various year-to-year growth rates are then put together to form a “chain” that can be used to compare the output of goods and services between any two dates.

This new chain-weighted measure of real GDP is better than the more traditional measure because it ensures that the prices used to compute real GDP are never far out of date. For most purposes, however, the differences are not important. It turns out that the two measures of real GDP are highly correlated with each other. The reason for this close association is that most relative prices change slowly over time. Thus, both measures of real GDP reflect the same thing: economy-wide changes in the production of goods and services.

FYI

Two Arithmetic Tricks for Working With Percentage Changes

For manipulating many relationships in economics, there is an arithmetic trick that is useful to know: *the percentage change of a product of two variables is approximately the sum of the percentage changes in each of the variables.*

To see how this trick works, consider an example. Let P denote the GDP deflator and Y denote real GDP. Nominal GDP is $P \times Y$. The trick states that

$$\begin{aligned} \text{Percentage Change in } (P \times Y) \\ \approx (\text{Percentage Change in } P) \\ + (\text{Percentage Change in } Y). \end{aligned}$$

For instance, suppose that in one year, real GDP is 100 and the GDP deflator is 2; the next year, real GDP is 103 and the GDP deflator is 2.1. We can calculate that real GDP rose by 3 percent and that the GDP deflator rose by 5 percent. Nominal GDP rose from 200 the first year to 216.3 the second year, an increase of 8.15 percent. Notice that the growth in nominal GDP

(8.15 percent) is approximately the sum of the growth in the GDP deflator (5 percent) and the growth in real GDP (3 percent).¹

A second arithmetic trick follows as a corollary to the first: *the percentage change of a ratio is approximately the percentage change in the numerator minus the percentage change in the denominator.* Again, consider an example. Let Y denote GDP and L denote the population, so that Y/L is GDP per person. The second trick states

$$\begin{aligned} \text{Percentage Change in } (Y/L) \\ \approx (\text{Percentage Change in } Y) \\ - (\text{Percentage Change in } L). \end{aligned}$$

For instance, suppose that in the first year, Y is 100,000 and L is 100, so Y/L is 1,000; in the second year, Y is 110,000 and L is 103, so Y/L is 1,068. Notice that the growth in GDP per person (6.8 percent) is approximately the growth in income (10 percent) minus the growth in population (3 percent).

The Components of Expenditure

Economists and policymakers care not only about the economy's total output of goods and services but also about the allocation of this output among alternative uses. The national income accounts divide GDP into four broad categories of spending:

- Consumption (C)
- Investment (I)
- Government purchases (G)
- Net exports (NX).

Thus, letting Y stand for GDP,

$$Y = C + I + G + NX.$$

¹ *Mathematical note:* The proof that this trick works begins with the chain rule from calculus:

$$d(PY) = Y dP + P dY.$$

Now divide both sides of this equation by PY to obtain

$$d(PY)/(PY) = dP/P + dY/Y.$$

Notice that all three terms in this equation are percentage changes.



What Is Investment?

Newcomers to macroeconomics are sometimes confused by how macroeconomists use familiar words in new and specific ways. One example is the term “investment.” The confusion arises because what looks like investment for an individual may not be investment for the economy as a whole. The general rule is that the economy’s investment does not include purchases that merely reallocate existing assets among different individuals. Investment, as macroeconomists use the term, creates new capital.

Let’s consider some examples. Suppose we observe these two events:

- Smith buys for himself a 100-year-old Victorian house.
- Jones builds for herself a brand-new contemporary house.

What is total investment here? Two houses, one house, or zero?

A macroeconomist seeing these two transactions counts only the Jones house as investment.

Smith’s transaction has not created new housing for the economy; it has merely reallocated existing housing. Smith’s purchase is investment for Smith, but it is disinvestment for the person selling the house. By contrast, Jones has added new housing to the economy; her new house is counted as investment.

Similarly, consider these two events:

- Gates buys \$5 million in IBM stock from Buffett on the New York Stock Exchange.
- General Motors sells \$10 million in stock to the public and uses the proceeds to build a new car factory.

Here, investment is \$10 million. In the first transaction, Gates is investing in IBM stock, and Buffett is disinvesting; there is no investment for the economy. By contrast, General Motors is using some of the economy’s output of goods and services to add to its stock of capital; hence, its new factory is counted as investment.

GDP is the sum of consumption, investment, government purchases, and net exports. Each dollar of GDP falls into one of these categories. This equation is an *identity*—an equation that must hold because of the way the variables are defined. It is called the **national income accounts identity**.

Consumption consists of the goods and services bought by households. It is divided into three subcategories: nondurable goods, durable goods, and services. Nondurable goods are goods that last only a short time, such as food and clothing. Durable goods are goods that last a long time, such as cars and TVs. Services include the work done for consumers by individuals and firms, such as haircuts and doctor visits.

Investment consists of goods bought for future use. Investment is also divided into three subcategories: business fixed investment, residential fixed investment, and inventory investment. Business fixed investment is the purchase of new plant and equipment by firms. Residential investment is the purchase of new housing by households and landlords. Inventory investment is the increase in firms’ inventories of goods (if inventories are falling, inventory investment is negative).

Government purchases are the goods and services bought by federal, state, and local governments. This category includes such items as military equipment, highways, and the services that government workers provide. It does not include

transfer payments to individuals, such as Social Security and welfare. Because transfer payments reallocate existing income and are not made in exchange for goods and services, they are not part of GDP.

The last category, **net exports**, takes into account trade with other countries. Net exports are the value of goods and services exported to other countries minus the value of goods and services that foreigners provide us. Net exports represent the net expenditure from abroad on our goods and services, which provides income for domestic producers.

CASE STUDY

GDP and Its Components

In 2000 the GDP of the United States totaled about \$10 trillion. This number is so large that it is almost impossible to comprehend. We can make it easier to understand by dividing it by the 2000 U.S. population of 275 million. In this way, we obtain GDP per person—the amount of expenditure for the average American—which equaled \$36,174 in 2000.

table 2-1

GDP and the Components of Expenditure: 2000

	Total (billions of dollars)	Per Person (dollars)
Gross Domestic Product	9,963.1	36,174
Consumption	6,757.3	24,534
Nondurable goods	2,010.0	7,298
Durable goods	820.3	2,978
Services	3,927.0	14,258
Investment	1,832.7	6,654
Nonresidential fixed investment	1,362.2	4,946
Residential fixed investment	416.0	1,510
Inventory investment	54.5	198
Government Purchases	1,743.7	6,331
Federal	595.2	2,161
Defense	377.0	1,369
Nondefense	218.2	792
State and local	1,148.6	4,170
Net Exports	−370.7	−1,346
Exports	1,097.3	3,984
Imports	1,468.0	5,330

Source: U.S. Department of Commerce.

How did this GDP get used? Table 2-1 shows that about two-thirds of it, or \$24,534 per person, was spent on consumption. Investment was \$6,654 per person. Government purchases were \$6,331 per person, \$1,369 of which was spent by the federal government on national defense.

The average American bought \$5,330 of goods imported from abroad and produced \$3,984 of goods that were exported to other countries. Because the average American imported more than he exported, net exports were negative. Furthermore, because the average American earned less from selling to foreigners than he spent on foreign goods, the difference must have been financed by taking out loans from foreigners (or, equivalently, by selling them some assets). Thus, the average American borrowed \$1,346 from abroad in 2000.

Other Measures of Income

The national income accounts include other measures of income that differ slightly in definition from GDP. It is important to be aware of the various measures, because economists and the press often refer to them.

To see how the alternative measures of income relate to one another, we start with GDP and add or subtract various quantities. To obtain *gross national product (GNP)*, we add receipts of factor income (wages, profit, and rent) from the rest of the world and subtract payments of factor income to the rest of the world:

$$\text{GNP} = \text{GDP} + \text{Factor Payments From Abroad} - \text{Factor Payments to Abroad}.$$

Whereas GDP measures the total income produced *domestically*, GNP measures the total income earned by *nationals* (residents of a nation). For instance, if a Japanese resident owns an apartment building in New York, the rental income he earns is part of U.S. GDP because it is earned in the United States. But because this rental income is a factor payment to abroad, it is not part of U.S. GNP. In the United States, factor payments from abroad and factor payments to abroad are similar in size—each representing about 3 percent of GDP—so GDP and GNP are quite close.

To obtain *net national product (NNP)*, we subtract the depreciation of capital—the amount of the economy’s stock of plants, equipment, and residential structures that wears out during the year:

$$\text{NNP} = \text{GNP} - \text{Depreciation}.$$

In the national income accounts, depreciation is called the *consumption of fixed capital*. It equals about 10 percent of GNP. Because the depreciation of capital is a cost of producing the output of the economy, subtracting depreciation shows the net result of economic activity.

The next adjustment in the national income accounts is for indirect business taxes, such as sales taxes. These taxes, which make up about 10 percent of NNP, place a wedge between the price that consumers pay for a good and the price

that firms receive. Because firms never receive this tax wedge, it is not part of their income. Once we subtract indirect business taxes from NNP, we obtain a measure called *national income*:

$$\text{National Income} = \text{NNP} - \text{Indirect Business Taxes.}$$

National income measures how much everyone in the economy has earned.

The national income accounts divide national income into five components, depending on the way the income is earned. The five categories, and the percentage of national income paid in each category, are

- *Compensation of employees* (70%). The wages and fringe benefits earned by workers.
- *Proprietors' income* (9%). The income of noncorporate businesses, such as small farms, mom-and-pop stores, and law partnerships.
- *Rental income* (2%). The income that landlords receive, including the imputed rent that homeowners “pay” to themselves, less expenses, such as depreciation.
- *Corporate profits* (12%). The income of corporations after payments to their workers and creditors.
- *Net interest* (7%). The interest domestic businesses pay minus the interest they receive, plus interest earned from foreigners.

A series of adjustments takes us from national income to *personal income*, the amount of income that households and noncorporate businesses receive. Three of these adjustments are most important. First, we reduce national income by the amount that corporations earn but do not pay out, either because the corporations are retaining earnings or because they are paying taxes to the government. This adjustment is made by subtracting corporate profits (which equals the sum of corporate taxes, dividends, and retained earnings) and adding back dividends. Second, we increase national income by the net amount the government pays out in transfer payments. This adjustment equals government transfers to individuals minus social insurance contributions paid to the government. Third, we adjust national income to include the interest that households earn rather than the interest that businesses pay. This adjustment is made by adding personal interest income and subtracting net interest. (The difference between personal interest and net interest arises in part from the interest on the government debt.) Thus, personal income is

$$\begin{aligned} \text{Personal Income} &= \text{National Income} \\ &\quad - \text{Corporate Profits} \\ &\quad - \text{Social Insurance Contributions} \\ &\quad - \text{Net Interest} \\ &\quad + \text{Dividends} \\ &\quad + \text{Government Transfers to Individuals} \\ &\quad + \text{Personal Interest Income.} \end{aligned}$$

Next, if we subtract personal tax payments and certain nontax payments to the government (such as parking tickets), we obtain *disposable personal income*:

$$\begin{aligned} &\text{Disposable Personal Income} \\ &= \text{Personal Income} - \text{Personal Tax and Nontax Payments.} \end{aligned}$$

We are interested in disposable personal income because it is the amount households and noncorporate businesses have available to spend after satisfying their tax obligations to the government.

CASE STUDY

The Seasonal Cycle and Seasonal Adjustment

Because real GDP and the other measures of income reflect how well the economy is performing, economists are interested in studying the quarter-to-quarter fluctuations in these variables. Yet when we start to do so, one fact leaps out: all these measures of income exhibit a regular seasonal pattern. The output of the economy rises during the year, reaching a peak in the fourth quarter (October, November, and December), and then falling in the first quarter (January, February, and March) of the next year. These regular seasonal changes are substantial. From the fourth quarter to the first quarter, real GDP falls on average about 8 percent.²

It is not surprising that real GDP follows a seasonal cycle. Some of these changes are attributable to changes in our ability to produce: for example, building homes is more difficult during the cold weather of winter than during other seasons. In addition, people have seasonal tastes: they have preferred times for such activities as vacations and holiday shopping.

When economists study fluctuations in real GDP and other economic variables, they often want to eliminate the portion of fluctuations caused by predictable seasonal changes. You will find that most of the economic statistics reported in the newspaper are *seasonally adjusted*. This means that the data have been adjusted to remove the regular seasonal fluctuations. (The precise statistical procedures used are too elaborate to bother with here, but in essence they involve subtracting those changes in income that are predictable simply from the change in season.) Therefore, when you observe a rise or fall in real GDP or any other data series, you must look beyond the seasonal cycle for the explanation.

² Robert B. Barsky and Jeffrey A. Miron, “The Seasonal Cycle and the Business Cycle,” *Journal of Political Economy* 97 (June 1989): 503–534.

2-2 Measuring the Cost of Living: The Consumer Price Index

A dollar today doesn't buy as much as it did 20 years ago. The cost of almost everything has gone up. This increase in the overall level of prices is called *inflation*, and it is one of the primary concerns of economists and policymakers. In later chapters we examine in detail the causes and effects of inflation. Here we discuss how economists measure changes in the cost of living.

The Price of a Basket of Goods

The most commonly used measure of the level of prices is the consumer price index (CPI). The Bureau of Labor Statistics, which is part of the U.S. Department of Labor, has the job of computing the CPI. It begins by collecting the prices of thousands of goods and services. Just as GDP turns the quantities of many goods and services into a single number measuring the value of production, the CPI turns the prices of many goods and services into a single index measuring the overall level of prices.

How should economists aggregate the many prices in the economy into a single index that reliably measures the price level? They could simply compute an average of all prices. Yet this approach would treat all goods and services equally. Because people buy more chicken than caviar, the price of chicken should have a greater weight in the CPI than the price of caviar. The Bureau of Labor Statistics weights different items by computing the price of a basket of goods and services purchased by a typical consumer. The CPI is the price of this basket of goods and services relative to the price of the same basket in some base year.

For example, suppose that the typical consumer buys 5 apples and 2 oranges every month. Then the basket of goods consists of 5 apples and 2 oranges, and the CPI is

$$\text{CPI} = \frac{(5 \times \text{Current Price of Apples}) + (2 \times \text{Current Price of Oranges})}{(5 \times 2002 \text{ Price of Apples}) + (2 \times 2002 \text{ Price of Oranges})}.$$

In this CPI, 2002 is the base year. The index tells us how much it costs now to buy 5 apples and 2 oranges relative to how much it cost to buy the same basket of fruit in 2002.

The consumer price index is the most closely watched index of prices, but it is not the only such index. Another is the producer price index, which measures the price of a typical basket of goods bought by firms rather than consumers. In addition to these overall price indices, the Bureau of Labor Statistics computes price indices for specific types of goods, such as food, housing, and energy.

The CPI Versus the GDP Deflator

Earlier in this chapter we saw another measure of prices—the implicit price deflator for GDP, which is the ratio of nominal GDP to real GDP. The GDP deflator and the CPI give somewhat different information about what’s happening to the overall level of prices in the economy. There are three key differences between the two measures.

The first difference is that the GDP deflator measures the prices of all goods and services produced, whereas the CPI measures the prices of only the goods and services bought by consumers. Thus, an increase in the price of goods bought by firms or the government will show up in the GDP deflator but not in the CPI.


The second difference is that the GDP deflator includes only those goods produced domestically. Imported goods are not part of GDP and do not show up in the GDP deflator. Hence, an increase in the price of a Toyota made in Japan and sold in this country affects the CPI, because the Toyota is bought by consumers, but it does not affect the GDP deflator.

The third and most subtle difference results from the way the two measures aggregate the many prices in the economy. The CPI assigns fixed weights to the prices of different goods, whereas the GDP deflator assigns changing weights. In other words, the CPI is computed using a fixed basket of goods, whereas the GDP deflator allows the basket of goods to change over time as the composition of GDP changes. The following example shows how these approaches differ. Suppose that major frosts destroy the nation’s orange crop. The quantity of oranges produced falls to zero, and the price of the few oranges that remain on grocers’ shelves is driven sky-high. Because oranges are no longer part of GDP, the increase in the price of oranges does not show up in the GDP deflator. But because the CPI is computed with a fixed basket of goods that includes oranges, the increase in the price of oranges causes a substantial rise in the CPI.

Economists call a price index with a fixed basket of goods a *Laspeyres index* and a price index with a changing basket a *Paasche index*. Economic theorists have studied the properties of these different types of price indices to determine which is a better measure of the cost of living. The answer, it turns out, is that neither is clearly superior. When prices of different goods are changing by different amounts, a Laspeyres (fixed basket) index tends to overstate the increase in the cost of living because it does not take into account that consumers have the opportunity to substitute less expensive goods for more expensive ones. By contrast, a Paasche (changing basket) index tends to understate the increase in the cost of living. Although it accounts for the substitution of alternative goods, it does not reflect the reduction in consumers’ welfare that may result from such substitutions.

The example of the destroyed orange crop shows the problems with Laspeyres and Paasche price indices. Because the CPI is a Laspeyres index, it overstates the impact of the increase in orange prices on consumers: by using a fixed basket of goods, it ignores consumers’ ability to substitute apples for oranges. By contrast,

figure 2-3



The GDP Deflator and the CPI This figure shows the percentage change in the GDP deflator and in the CPI for every year since 1948. Although these two measures of prices diverge at times, they usually tell the same story about how quickly prices are rising. Both the CPI and the GDP deflator show that prices rose slowly in most of the 1950s and 1960s, that they rose much more quickly in the 1970s, and that they rose slowly again in the 1980s and 1990s.

Source: U.S. Department of Commerce, U.S. Department of Labor.

because the GDP deflator is a Paasche index, it understates the impact on consumers: the GDP deflator shows no rise in prices, yet surely the higher price of oranges makes consumers worse off.

Luckily, the difference between the GDP deflator and the CPI is usually not large in practice. Figure 2-3 shows the percentage change in the GDP deflator and the percentage change in the CPI for each year since 1948. Both measures usually tell the same story about how quickly prices are rising.

CASE STUDY

Does the CPI Overstate Inflation?

The consumer price index is a closely watched measure of inflation. Policymakers in the Federal Reserve monitor the CPI when choosing monetary policy. In addition, many laws and private contracts have cost-of-living allowances, called COLAs, which use the CPI to adjust for changes in the price level. For instance,

Social Security benefits are adjusted automatically every year so that inflation will not erode the living standard of the elderly.

Because so much depends on the CPI, it is important to ensure that this measure of the price level is accurate. Many economists believe that, for a number of reasons, the CPI tends to overstate inflation.

One problem is the substitution bias we have already discussed. Because the CPI measures the price of a fixed basket of goods, it does not reflect the ability of consumers to substitute toward goods whose relative prices have fallen. Thus, when relative prices change, the true cost of living rises less rapidly than the CPI.

A second problem is the introduction of new goods. When a new good is introduced into the marketplace, consumers are better off, because they have more products from which to choose. In effect, the introduction of new goods increases the real value of the dollar. Yet this increase in the purchasing power of the dollar is not reflected in a lower CPI.

A third problem is unmeasured changes in quality. When a firm changes the quality of a good it sells, not all of the good's price change reflects a change in the cost of living. The Bureau of Economic Analysis does its best to account for changes in the quality of goods over time. For example, if Ford increases the horsepower of a particular car model from one year to the next, the CPI will reflect the change: the quality-adjusted price of the car will not rise as fast as the unadjusted price. Yet many changes in quality, such as comfort or safety, are hard to measure. If unmeasured quality improvement (rather than unmeasured quality deterioration) is typical, then the measured CPI rises faster than it should.

Because of these measurement problems, some economists have suggested revising laws to reduce the degree of indexation. For example, Social Security benefits could be indexed to CPI inflation minus 1 percent. Such a change would provide a rough way of offsetting these measurement problems. At the same time, it would automatically slow the growth in government spending.

In 1995, the Senate Finance Committee appointed a panel of five noted economists—Michael Boskin, Ellen Dulberger, Robert Gordon, Zvi Griliches, and Dale Jorgenson—to study the magnitude of the measurement error in the CPI. The panel concluded that the CPI was biased upward by 0.8 to 1.6 percentage points per year, with their “best estimate” being 1.1 percentage points. This report led to some changes in the way the CPI is calculated, so the bias is now thought to be slightly under 1 percentage point. The CPI still overstates inflation, but not by as much as it once did.³

³ For further discussion of these issues, see Matthew Shapiro and David Wilcox, “Mismeasurement in the Consumer Price Index: An Evaluation,” *NBER Macroeconomics Annual*, 1996, and the symposium on “Measuring the CPI” in the Winter 1998 issue of *The Journal of Economic Perspectives*.

2-3 Measuring Joblessness: The Unemployment Rate

One aspect of economic performance is how well an economy uses its resources. Because an economy's workers are its chief resource, keeping workers employed is a paramount concern of economic policymakers. The unemployment rate is the statistic that measures the percentage of those people wanting to work who do not have jobs.



Drawing M. Stevens, © 1980 The New Yorker Magazine, Inc.

“Well, so long Eddie, the recession’s over.”

Every month the U.S. Bureau of Labor Statistics computes the unemployment rate and many other statistics that economists and policymakers use to monitor developments in the labor market. These statistics come from a survey of about 60,000 households. Based on the responses to survey questions, each adult (16 years and older) in each household is placed into one of three categories: employed, unemployed, or not in the labor force. A person is employed if he or she spent some of the previous week working at a paid job. A person is unemployed if he or she is not employed and has been looking for a job or is on temporary layoff. A person who fits into neither of the first two categories, such as a full-time student or retiree, is not in the labor force. A person who wants a job but has given up looking—a *discouraged worker*—is counted as not being in the labor force.

The **labor force** is defined as the sum of the employed and unemployed, and the unemployment rate is defined as the percentage of the labor force that is unemployed. That is,

$$\text{Labor Force} = \text{Number of Employed} + \text{Number of Unemployed},$$

and


$$\text{Unemployment Rate} = \frac{\text{Number of Unemployed}}{\text{Labor Force}} \times 100.$$

A related statistic is the **labor-force participation rate**, the percentage of the adult population that is in the labor force:

$$\text{Labor-Force Participation Rate} = \frac{\text{Labor Force}}{\text{Adult Population}} \times 100.$$

The Bureau of Labor Statistics computes these statistics for the overall population and for groups within the population: men and women, whites and blacks, teenagers and prime-age workers.

figure 2-4



The Three Groups of the Population

When the Bureau of Labor Statistics surveys the population, it places all adults into one of three categories: employed, unemployed, or not in the labor force. This figure shows the number of people in each category in 2000.

Source: U.S. Department of Labor.

Figure 2-4 shows the breakdown of the population into the three categories for 2000. The statistics broke down as follows:

$$\text{Labor Force} = 135.2 + 5.7 = 140.9 \text{ million.}$$

$$\text{Unemployment Rate} = (5.7/140.9) \times 100 = 4.0\%.$$

$$\text{Labor-Force Participation Rate} = (140.9/209.7) \times 100 = 67.2\%.$$

Hence, about two-thirds of the adult population was in the labor force, and about 4 percent of those in the labor force did not have a job.


CASE STUDY

Unemployment, GDP, and Okun's Law

What relationship should we expect to find between unemployment and real GDP? Because employed workers help to produce goods and services and unemployed workers do not, increases in the unemployment rate should be associated with decreases in real GDP. This negative relationship between unemployment and GDP is called **Okun's law**, after Arthur Okun, the economist who first studied it.⁴

⁴ Arthur M. Okun, "Potential GNP: Its Measurement and Significance," in *Proceedings of the Business and Economics Statistics Section, American Statistical Association* (Washington, DC: American Statistical Association, 1962), 98–103; reprinted in Arthur M. Okun, *Economics for Policymaking* (Cambridge, MA: MIT Press, 1983), 145–158.

figure 2-5



Okun's Law This figure is a scatterplot of the change in the unemployment rate on the horizontal axis and the percentage change in real GDP on the vertical axis, using data on the U.S. economy. Each point represents one year. The negative correlation between these variables shows that increases in unemployment tend to be associated with lower-than-normal growth in real GDP.

Source: U.S. Department of Commerce, U.S. Department of Labor.

Figure 2-5 uses annual data for the United States to illustrate Okun's law. This figure is a scatterplot—a scatter of points where each point represents one observation (in this case, the data for a particular year). The horizontal axis represents the change in the unemployment rate from the previous year, and the vertical axis represents the percentage change in GDP. This figure shows clearly that year-to-year changes in the unemployment rate are closely associated with year-to-year changes in real GDP.

We can be more precise about the magnitude of the Okun's law relationship. The line drawn through the scatter of points (estimated with a statistical procedure called ordinary least squares) tells us that

$$\begin{aligned} \text{Percentage Change in Real GDP} \\ = 3\% - 2 \times \text{Change in the Unemployment Rate.} \end{aligned}$$

If the unemployment rate remains the same, real GDP grows by about 3 percent; this normal growth in the production of goods and services is a result of growth in the labor force, capital accumulation, and technological progress. In addition, for every percentage point the unemployment rate rises, real GDP growth typi-

cally falls by 2 percent. Hence, if the unemployment rate rises from 6 to 8 percent, then real GDP growth would be

$$\begin{aligned}\text{Percentage Change in Real GDP} &= 3\% - 2 \times (8\% - 6\%) \\ &= -1\%.\end{aligned}$$

In this case, Okun's law says that GDP would fall by 1 percent, indicating that the economy is in a recession.

2-4 Conclusion: From Economic Statistics to Economic Models

The three statistics discussed in this chapter—gross domestic product, the consumer price index, and the unemployment rate—quantify the performance of the economy. Public and private decisionmakers use these statistics to monitor changes in the economy and to formulate appropriate policies. Economists use these statistics to develop and test theories about how the economy works.

In the chapters that follow, we examine some of these theories. That is, we build models that explain how these variables are determined and how economic policy affects them. Having learned how to measure economic performance, we are now ready to learn how to explain it.

Summary

1. Gross domestic product (GDP) measures both the income of everyone in the economy and the total expenditure on the economy's output of goods and services.
2. Nominal GDP values goods and services at current prices. Real GDP values goods and services at constant prices. Real GDP rises only when the amount of goods and services has increased, whereas nominal GDP can rise either because output has increased or because prices have increased.
3. GDP is the sum of four categories of expenditure: consumption, investment, government purchases, and net exports.
4. The consumer price index (CPI) measures the price of a fixed basket of goods and services purchased by a typical consumer. Like the GDP deflator, which is the ratio of nominal GDP to real GDP, the CPI measures the overall level of prices.
5. The unemployment rate shows what fraction of those who would like to work do not have a job. When the unemployment rate rises, real GDP typically grows slower than its normal rate and may even fall.

KEY CONCEPTS

Gross domestic product (GDP)	Imputed value	Government purchases
Consumer price index (CPI)	Nominal versus real GDP	Net exports
Unemployment rate	GDP deflator	Labor force
National income accounting	National income accounts identity	Labor-force participation rate
Stocks and flows	Consumption	Okun's law
Value added	Investment	

QUESTIONS FOR REVIEW

1. List the two things that GDP measures. How can GDP measure two things at once?
2. What does the consumer price index measure?
3. List the three categories used by the Bureau of Labor Statistics to classify everyone in the economy. How does the bureau compute the unemployment rate?
4. Explain Okun's law.

PROBLEMS AND APPLICATIONS

1. Look at the newspapers for the past few days. What new economic statistics have been released? How do you interpret these statistics?
2. A farmer grows a bushel of wheat and sells it to a miller for \$1.00. The miller turns the wheat into flour and then sells the flour to a baker for \$3.00. The baker uses the flour to make bread and sells the bread to an engineer for \$6.00. The engineer eats the bread. What is the value added by each person? What is GDP?
3. Suppose a woman marries her butler. After they are married, her husband continues to wait on her as before, and she continues to support him as before (but as a husband rather than as an employee). How does the marriage affect GDP? How should it affect GDP?
4. Place each of the following transactions in one of the four components of expenditure: consumption, investment, government purchases, and net exports.
 - a. Boeing sells an airplane to the Air Force.
 - b. Boeing sells an airplane to American Airlines.
 - c. Boeing sells an airplane to Air France.
 - d. Boeing sells an airplane to Amelia Earhart.
 - e. Boeing builds an airplane to be sold next year.
5. Find data on GDP and its components, and compute the percentage of GDP for the following components for 1950, 1975, and 2000.
 - a. Personal consumption expenditures
 - b. Gross private domestic investment
 - c. Government purchases
 - d. Net exports
 - e. National defense purchases
 - f. State and local purchases
 - g. Imports

Do you see any stable relationships in the data? Do you see any trends? (*Hint:* A good place to look for data is the statistical appendices of the *Economic Report of the President*, which is written each year by the Council of Economic Advisers. Alternatively, you can go to www.bea.doc.gov, which is the Web site of the Bureau of Economic Analysis.)
6. Consider an economy that produces and consumes bread and automobiles. In the following table are data for two different years.

	Year 2000	Year 2010
Price of an automobile	\$50,000	\$60,000
Price of a loaf of bread	\$10	\$20
Number of automobiles produced	100	120
Number of loaves of bread produced	500,000	400,000

- a. Using the year 2000 as the base year, compute the following statistics for each year: nominal GDP, real GDP, the implicit price deflator for GDP, and a fixed-weight price index such as the CPI.
 - b. How much have prices risen between year 2000 and year 2010? Compare the answers given by the Laspeyres and Paasche price indices. Explain the difference.
 - c. Suppose you are a senator writing a bill to index Social Security and federal pensions. That is, your bill will adjust these benefits to offset changes in the cost of living. Will you use the GDP deflator or the CPI? Why?
7. Abby consumes only apples. In year 1, red apples cost \$1 each, green apples cost \$2 each, and Abby buys 10 red apples. In year 2, red apples cost \$2, green apples cost \$1, and Abby buys 10 green apples.
- a. Compute a consumer price index for apples for each year. Assume that year 1 is the base year in which the consumer basket is fixed. How does your index change from year 1 to year 2?
 - b. Compute Abby's nominal spending on apples in each year. How does it change from year 1 to year 2?
 - c. Using year 1 as the base year, compute Abby's real spending on apples in each year. How does it change from year 1 to year 2?
 - d. Defining the implicit price deflator as nominal spending divided by real spending, compute the deflator for each year. How does the deflator change from year 1 to year 2?
- e. Suppose that Abby is equally happy eating red or green apples. How much has the true cost of living increased for Abby? Compare this answer to your answers to parts (a) and (d). What does this example tell you about Laspeyres and Paasche price indexes?
8. Consider how each of the following events is likely to affect real GDP. Do you think the change in real GDP reflects a similar change in economic well-being?
- a. A hurricane in Florida forces Disney World to shut down for a month.
 - b. The discovery of a new, easy-to-grow strain of wheat increases farm harvests.
 - c. Increased hostility between unions and management sparks a rash of strikes.
 - d. Firms throughout the economy experience falling demand, causing them to lay off workers.
 - e. Congress passes new environmental laws that prohibit firms from using production methods that emit large quantities of pollution.
 - f. More high-school students drop out of school to take jobs mowing lawns.
 - g. Fathers around the country reduce their work-weeks to spend more time with their children.
9. In a speech that Senator Robert Kennedy gave when he was running for president in 1968, he said the following about GDP:
- [It] does not allow for the health of our children, the quality of their education, or the joy of their play. It does not include the beauty of our poetry or the strength of our marriages, the intelligence of our public debate or the integrity of our public officials. It measures neither our courage, nor our wisdom, nor our devotion to our country. It measures everything, in short, except that which makes life worthwhile, and it can tell us everything about America except why we are proud that we are Americans.
- Was Robert Kennedy right? If so, why do we care about GDP?

part II

**Classical Theory:
The Economy in
the Long Run**

National Income: Where It Comes From and Where It Goes

A large income is the best recipe for happiness I ever heard of.

— Jane Austen

The most important macroeconomic variable is gross domestic product (GDP). As we have seen, GDP measures both a nation's total output of goods and services and its total income. To appreciate the significance of GDP, one need only take a quick look at international data: compared with their poorer counterparts, nations with a high level of GDP per person have everything from better childhood nutrition to more televisions per household. A large GDP does not ensure that all of a nation's citizens are happy, but it may be the best recipe for happiness that macroeconomists have to offer.


This chapter addresses four groups of questions about the sources and uses of a nation's GDP:

- How much do the firms in the economy produce? What determines a nation's total income?
- Who gets the income from production? How much goes to compensate workers, and how much goes to compensate owners of capital?
- Who buys the output of the economy? How much do households purchase for consumption, how much do households and firms purchase for investment, and how much does the government buy for public purposes?
- What equilibrates the demand for and supply of goods and services? What ensures that desired spending on consumption, investment, and government purchases equals the level of production?

To answer these questions, we must examine how the various parts of the economy interact.

A good place to start is the circular flow diagram. In Chapter 2 we traced the circular flow of dollars in a hypothetical economy that produced one product, bread, from labor services. Figure 3-1 more accurately reflects how real economies function. It shows the linkages among the economic actors—households, firms,

figure 3-1



The Circular Flow of Dollars Through the Economy This figure is a more realistic version of the circular flow diagram found in Chapter 2. Each yellow box represents an economic actor—households, firms, and the government. Each blue box represents a type of market—the markets for goods and services, the markets for the factors of production, and financial markets. The green arrows show the flow of dollars among the economic actors through the three types of markets.

and the government—and how dollars flow among them through the various markets in the economy.

Let's look at the flow of dollars from the viewpoints of these economic actors. Households receive income and use it to pay taxes to the government, to consume goods and services, and to save through the financial markets. Firms receive revenue from the sale of goods and services and use it to pay for the factors of production. Both households and firms borrow in financial markets to buy investment goods, such as houses and factories. The government receives revenue from taxes and uses it to pay for government purchases. Any excess of tax revenue over government spending is called public saving, which can be either positive (a budget surplus) or negative (a budget deficit).

In this chapter we develop a basic classical model to explain the economic interactions depicted in Figure 3-1. We begin with firms and look at what

determines their level of production (and, thus, the level of national income). Then we examine how the markets for the factors of production distribute this income to households. Next, we consider how much of this income households consume and how much they save. In addition to discussing the demand for goods and services arising from the consumption of households, we discuss the demand arising from investment and government purchases. Finally, we come full circle and examine how the demand for goods and services (the sum of consumption, investment, and government purchases) and the supply of goods and services (the level of production) are brought into balance.

3-1 What Determines the Total Production of Goods and Services?

An economy's output of goods and services—its GDP—depends on (1) its quantity of inputs, called the factors of production, and (2) its ability to turn inputs into output, as represented by the production function. We discuss each of these in turn.

The Factors of Production

Factors of production are the inputs used to produce goods and services. The two most important factors of production are capital and labor. Capital is the set of tools that workers use: the construction worker's crane, the accountant's calculator, and this author's personal computer. Labor is the time people spend working. We use the symbol K to denote the amount of capital and the symbol L to denote the amount of labor.

In this chapter we take the economy's factors of production as given. In other words, we assume that the economy has a fixed amount of capital and a fixed amount of labor. We write

$$\begin{aligned}K &= \bar{K}. \\L &= \bar{L}.\end{aligned}$$

The overbar means that each variable is fixed at some level. In Chapter 7 we examine what happens when the factors of production change over time, as they do in the real world. For now, to keep our analysis simple, we assume fixed amounts of capital and labor.

We also assume here that the factors of production are fully utilized—that is, that no resources are wasted. Again, in the real world, part of the labor force is unemployed, and some capital lies idle. In Chapter 6 we examine the reasons for unemployment, but for now we assume that capital and labor are fully employed.

The Production Function

The available production technology determines how much output is produced from given amounts of capital and labor. Economists express the available technology using a **production function**. Letting Y denote the amount of output, we write the production function as

$$Y = F(K, L).$$

This equation states that output is a function of the amount of capital and the amount of labor.

The production function reflects the available technology for turning capital and labor into output. If someone invents a better way to produce a good, the result is more output from the same amounts of capital and labor. Thus, technological change alters the production function.

Many production functions have a property called **constant returns to scale**. A production function has constant returns to scale if an increase of an equal percentage in all factors of production causes an increase in output of the same percentage. If the production function has constant returns to scale, then we get 10 percent more output when we increase both capital and labor by 10 percent. Mathematically, a production function has constant returns to scale if

$$zY = F(zK, zL)$$

for any positive number z . This equation says that if we multiply both the amount of capital and the amount of labor by some number z , output is also multiplied by z . In the next section we see that the assumption of constant returns to scale has an important implication for how the income from production is distributed.

As an example of a production function, consider production at a bakery. The kitchen and its equipment are the bakery's capital, the workers hired to make the bread are its labor, and the loaves of bread are its output. The bakery's production function shows that the number of loaves produced depends on the amount of equipment and the number of workers. If the production function has constant returns to scale, then doubling the amount of equipment and the number of workers doubles the amount of bread produced.

The Supply of Goods and Services

We can now see that the factors of production and the production function together determine the quantity of goods and services supplied, which in turn equals the economy's output. To express this mathematically, we write

$$\begin{aligned} Y &= F(\bar{K}, \bar{L}) \\ &= \bar{Y}. \end{aligned}$$

In this chapter, because we assume that the supplies of capital and labor and the technology are fixed, output is also fixed (at a level denoted here as $\bar{Y}$). When we

discuss economic growth in Chapters 7 and 8, we will examine how increases in capital and labor and improvements in the production technology lead to growth in the economy's output.

3-2 How Is National Income Distributed to the Factors of Production?

As we discussed in Chapter 2, the total output of an economy equals its total income. Because the factors of production and the production function together determine the total output of goods and services, they also determine national income. The circular flow diagram in Figure 3-1 shows that this national income flows from firms to households through the markets for the factors of production.

In this section we continue developing our model of the economy by discussing how these factor markets work. Economists have long studied factor markets to understand the distribution of income. (For example, Karl Marx, the noted nineteenth-century economist, spent much time trying to explain the incomes of capital and labor. The political philosophy of communism was in part based on Marx's now-discredited theory.) Here we examine the modern theory of how national income is divided among the factors of production. This theory, called the *neoclassical theory of distribution*, is accepted by most economists today.

Factor Prices


The distribution of national income is determined by factor prices. **Factor prices** are the amounts paid to the factors of production—the wage workers earn and the rent the owners of capital collect. As Figure 3-2 illustrates, the price each factor of production receives for its services is in turn determined by the supply and demand for that factor. Because we have assumed that the economy's factors of production are fixed, the factor supply curve in Figure 3-2 is vertical. The intersection of the downward-sloping factor demand curve and the vertical supply curve determines the equilibrium factor price.

To understand factor prices and the distribution of income, we must examine the demand for the factors of production. Because factor demand arises from the thousands of firms that use capital and labor, we now look at the decisions faced by a typical firm about how much of these factors to employ.

The Decisions Facing the Competitive Firm

The simplest assumption to make about a typical firm is that it is **competitive**. A competitive firm is small relative to the markets in which it trades, so it has little influence on market prices. For example, our firm produces a good and sells it at the market price. Because many firms produce this good, our firm can sell as

figure 3-2



How a Factor of Production Is Compensated The price paid to any factor of production depends on the supply and demand for that factor's services. Because we have assumed that supply is fixed, the supply curve is vertical. The demand curve is downward sloping. The intersection of supply and demand determines the equilibrium factor price.

much as it wants without causing the price of the good to fall, or it can stop selling altogether without causing the price of the good to rise. Similarly, our firm cannot influence the wages of the workers it employs because many other local firms also employ workers. The firm has no reason to pay more than the market wage, and if it tried to pay less, its workers would take jobs elsewhere. Therefore, the competitive firm takes the prices of its output and its inputs as given.

To make its product, the firm needs two factors of production, capital and labor. As we did for the aggregate economy, we represent the firm's production technology by the production function

$$Y = F(K, L),$$

where Y is the number of units produced (the firm's output), K the number of machines used (the amount of capital), and L the number of hours worked by the firm's employees (the amount of labor). The firm produces more output if it has more machines or if its employees work more hours.

The firm sells its output at a price P , hires workers at a wage W , and rents capital at a rate R . Notice that when we speak of firms renting capital, we are assuming that households own the economy's stock of capital. In this analysis, households rent out their capital, just as they sell their labor. The firm obtains both factors of production from the households that own them.¹

The goal of the firm is to maximize profit. *Profit* is revenue minus costs—it is what the owners of the firm keep after paying for the costs of production. Revenue equals $P \times Y$, the selling price of the good P multiplied by the amount of

¹ This is a simplification. In the real world, the ownership of capital is indirect because firms own capital and households own the firms. That is, real firms have two functions: owning capital and producing output. To help us understand how the factors of production are compensated, however, we assume that firms only produce output and that households own capital directly.

the good the firm produces Y . Costs include both labor costs and capital costs. Labor costs equal $W \times L$, the wage W times the amount of labor L . Capital costs equal $R \times K$, the rental price of capital R times the amount of capital K . We can write

$$\begin{aligned}\text{Profit} &= \text{Revenue} - \text{Labor Costs} - \text{Capital Costs} \\ &= PY - WL - RK.\end{aligned}$$

To see how profit depends on the factors of production, we use the production function $Y = F(K, L)$ to substitute for Y to obtain

$$\text{Profit} = PF(K, L) - WL - RK.$$

This equation shows that profit depends on the product price P , the factor prices W and R , and the factor quantities L and K . The competitive firm takes the product price and the factor prices as given and chooses the amounts of labor and capital that maximize profit.

The Firm's Demand for Factors

We now know that our firm will hire labor and rent capital in the quantities that maximize profit. But what are those profit-maximizing quantities? To answer this question, we first consider the quantity of labor and then the quantity of capital.

The Marginal Product of Labor The more labor the firm employs, the more output it produces. The **marginal product of labor (MPL)** is the extra amount of output the firm gets from one extra unit of labor, holding the amount of capital fixed. We can express this using the production function:


$$MPL = F(K, L + 1) - F(K, L).$$

The first term on the right-hand side is the amount of output produced with K units of capital and $L + 1$ units of labor; the second term is the amount of output produced with K units of capital and L units of labor. This equation states that the marginal product of labor is the difference between the amount of output produced with $L + 1$ units of labor and the amount produced with only L units of labor.

Most production functions have the property of **diminishing marginal product**: holding the amount of capital fixed, the marginal product of labor decreases as the amount of labor increases. Consider again the production of bread at a bakery. As a bakery hires more labor, it produces more bread. The MPL is the amount of extra bread produced when an extra unit of labor is hired. As more labor is added to a fixed amount of capital, however, the MPL falls. Fewer additional loaves are produced because workers are less productive when the kitchen is more crowded. In other words, holding the size of the kitchen fixed, each additional worker adds fewer loaves of bread to the bakery's output.

Figure 3-3 graphs the production function. It illustrates what happens to the amount of output when we hold the amount of capital constant and vary the

figure 3-3



The Production Function This curve shows how output depends on labor input, holding the amount of capital constant. The marginal product of labor MPL is the change in output when the labor input is increased by 1 unit. As the amount of labor increases, the production function becomes flatter, indicating diminishing marginal product.

amount of labor. This figure shows that the marginal product of labor is the slope of the production function. As the amount of labor increases, the production function becomes flatter, indicating diminishing marginal product.

From the Marginal Product of Labor to Labor Demand When the competitive, profit-maximizing firm is deciding whether to hire an additional unit of labor, it considers how that decision would affect profits. It therefore compares the extra revenue from the increased production that results from the added labor to the extra cost of higher spending on wages. The increase in revenue from an additional unit of labor depends on two variables: the marginal product of labor and the price of the output. Because an extra unit of labor produces MPL units of output and each unit of output sells for P dollars, the extra revenue is $P \times MPL$. The extra cost of hiring one more unit of labor is the wage W . Thus, the change in profit from hiring an additional unit of labor is

$$\begin{aligned}\Delta \text{ Profit} &= \Delta \text{ Revenue} - \Delta \text{ Cost} \\ &= (P \times MPL) - W.\end{aligned}$$

The symbol Δ (called *delta*) denotes the change in a variable.

We can now answer the question we asked at the beginning of this section: How much labor does the firm hire? The firm's manager knows that if the extra revenue $P \times MPL$ exceeds the wage W , an extra unit of labor increases profit. Therefore, the manager continues to hire labor until the next unit would no longer be profitable—that is, until the MPL falls to the point where the extra revenue equals the wage. The firm's demand for labor is determined by

$$P \times MPL = W.$$

We can also write this as


$$MPL = W/P.$$

W/P is the **real wage**—the payment to labor measured in units of output rather than in dollars. To maximize profit, the firm hires up to the point at which the marginal product of labor equals the real wage.

For example, again consider a bakery. Suppose the price of bread P is \$2 per loaf, and a worker earns a wage W of \$20 per hour. The real wage W/P is 10 loaves per hour. In this example, the firm keeps hiring workers as long as each additional worker would produce at least 10 loaves per hour. When the MPL falls to 10 loaves per hour or less, hiring additional workers is no longer profitable.

Figure 3-4 shows how the marginal product of labor depends on the amount of labor employed (holding the firm's capital stock constant). That is, this figure graphs the MPL schedule. Because the MPL diminishes as the amount of labor increases, this curve slopes downward. For any given real wage, the firm hires up to the point at which the MPL equals the real wage. Hence, the MPL schedule is also the firm's labor demand curve.

figure 3-4



The Marginal Product of Labor Schedule The marginal product of labor MPL depends on the amount of labor. The MPL curve slopes downward because the MPL declines as L increases. The firm hires labor up to the point where the real wage W/P equals the MPL . Hence, this schedule is also the firm's labor demand curve.

The Marginal Product of Capital and Capital Demand The firm decides how much capital to rent in the same way it decides how much labor to hire. The **marginal product of capital (MPK)** is the amount of extra output the firm gets from an extra unit of capital, holding the amount of labor constant:

$$MPK = F(K + 1, L) - F(K, L).$$

Thus, the marginal product of capital is the difference between the amount of output produced with $K + 1$ units of capital and that produced with only K units of capital. Like labor, capital is subject to diminishing marginal product.

The increase in profit from renting an additional machine is the extra revenue from selling the output of that machine minus the machine's rental price:

$$\begin{aligned}\Delta \text{Profit} &= \Delta \text{Revenue} - \Delta \text{Cost} \\ &= (P \times MPK) - R.\end{aligned}$$

To maximize profit, the firm continues to rent more capital until the MPK falls to equal the real rental price:

$$MPK = R/P.$$

The **real rental price of capital** is the rental price measured in units of goods rather than in dollars.

To sum up, the competitive, profit-maximizing firm follows a simple rule about how much labor to hire and how much capital to rent. *The firm demands each factor of production until that factor's marginal product falls to equal its real factor price.*

The Division of National Income

Having analyzed how a firm decides how much of each factor to employ, we can now explain how the markets for the factors of production distribute the economy's total income. If all firms in the economy are competitive and profit maximizing, then each factor of production is paid its marginal contribution to the production process. The real wage paid to each worker equals the MPL , and the real rental price paid to each owner of capital equals the MPK . The total real wages paid to labor are therefore $MPL \times L$, and the total real return paid to capital owners is $MPK \times K$.

The income that remains after the firms have paid the factors of production is the **economic profit** of the owners of the firms. Real economic profit is

$$\text{Economic Profit} = Y - (MPL \times L) - (MPK \times K).$$

Because we want to examine the distribution of national income, we rearrange the terms as follows:

$$Y = (MPL \times L) + (MPK \times K) + \text{Economic Profit}.$$

Total income is divided among the return to labor, the return to capital, and economic profit.

How large is economic profit? The answer is surprising: if the production function has the property of constant returns to scale, as is often thought to be the case, then economic profit must be zero. That is, nothing is left after the factors of production are paid. This conclusion follows from a famous mathematical result called *Euler's theorem*,² which states that if the production function has constant returns to scale, then

$$F(K, L) = (MPK \times K) + (MPL \times L).$$

If each factor of production is paid its marginal product, then the sum of these factor payments equals total output. In other words, constant returns to scale, profit maximization, and competition together imply that economic profit is zero.

If economic profit is zero, how can we explain the existence of “profit” in the economy? The answer is that the term “profit” as normally used is different from economic profit. We have been assuming that there are three types of agents: workers, owners of capital, and owners of firms. Total income is divided among wages, return to capital, and economic profit. In the real world, however, most firms own rather than rent the capital they use. Because firm owners and capital owners are the same people, economic profit and the return to capital are often lumped together. If we call this alternative definition **accounting profit**, we can say that

$$\text{Accounting Profit} = \text{Economic Profit} + (MPK \times K).$$

Under our assumptions—constant returns to scale, profit maximization, and competition—economic profit is zero. If these assumptions approximately describe the world, then the “profit” in the national income accounts must be mostly the return to capital.

We can now answer the question posed at the beginning of this chapter about how the income of the economy is distributed from firms to households. Each factor of production is paid its marginal product, and these factor payments exhaust total output. *Total output is divided between the payments to capital and the payments to labor, depending on their marginal productivities.*

CASE STUDY

The Black Death and Factor Prices

As we have just learned, in the neoclassical theory of distribution, factor prices equal the marginal products of the factors of production. Because the marginal products depend on the quantities of the factors, a change in the quantity of any one factor alters the marginal products of all the factors. Therefore, a change in the supply of a factor alters equilibrium factor prices.

Fourteenth-century Europe provides a vivid example of how factor quantities affect factor prices. The outbreak of the bubonic plague—the Black Death—in

² *Mathematical note:* To prove Euler's theorem, begin with the definition of constant returns to scale: $zY = F(zK, zL)$. Now differentiate with respect to z and then evaluate at $z = 1$.

1348 reduced the population of Europe by about one-third within a few years. Because the marginal product of labor increases as the amount of labor falls, this massive reduction in the labor force raised the marginal product of labor. (The economy moved to the left along the curves in Figures 3-3 and 3-4.) Real wages did increase substantially during the plague years—doubling, by some estimates. The peasants who were fortunate enough to survive the plague enjoyed economic prosperity.

The reduction in the labor force caused by the plague also affected the return to land, the other major factor of production in medieval Europe. With fewer workers available to farm the land, an additional unit of land produced less additional output. This fall in the marginal product of land led to a decline in real rents of 50 percent or more. Thus, while the peasant classes prospered, the landed classes suffered reduced incomes.³

3-3 What Determines the Demand for Goods and Services?

We have seen what determines the level of production and how the income from production is distributed to workers and owners of capital. We now continue our tour of the circular flow diagram, Figure 3-1, and examine how the output from production is used.

In Chapter 2 we identified the four components of GDP:

- Consumption (C)
- Investment (I)
- Government purchases (G)
- Net exports (NX).

The circular flow diagram contains only the first three components. For now, to simplify the analysis, we assume a *closed economy*—a country that does not trade with other countries. Thus, net exports are always zero. (We examine the macroeconomics of *open economies* in Chapter 5.)

A closed economy has three uses for the goods and services it produces. These three components of GDP are expressed in the national income accounts identity:

$$Y = C + I + G.$$

Households consume some of the economy's output; firms and households use some of the output for investment; and the government buys some of the output for public purposes. We want to see how GDP is allocated among these three uses.

³ Carlo M. Cipolla, *Before the Industrial Revolution: European Society and Economy, 1000–1700*, 2d ed. (New York: Norton, 1980), 200–202.

Consumption

When we eat food, wear clothing, or go to a movie, we are consuming some of the output of the economy. All forms of consumption together make up two-thirds of GDP. Because consumption is so large, macroeconomists have devoted much energy to studying how households decide how much to consume. Chapter 16 examines this work in detail. Here we consider the simplest story of consumer behavior.

Households receive income from their labor and their ownership of capital, pay taxes to the government, and then decide how much of their after-tax income to consume and how much to save. As we discussed in Section 3-2, the income that households receive equals the output of the economy Y . The government then taxes households an amount T . (Although the government imposes many kinds of taxes, such as personal and corporate income taxes and sales taxes, for our purposes we can lump all these taxes together.) We define income after the payment of all taxes, $Y - T$, as **disposable income**. Households divide their disposable income between consumption and saving.

We assume that the level of consumption depends directly on the level of disposable income. The higher the disposable income, the greater the consumption. Thus,

$$C = C(Y - T).$$

This equation states that consumption is a function of disposable income. The relationship between consumption and disposable income is called the **consumption function**.

The **marginal propensity to consume (MPC)** is the amount by which consumption changes when disposable income increases by one dollar. The *MPC* is between zero and one: an extra dollar of income increases consumption, but by less than one dollar. Thus, if households obtain an extra dollar of income, they save a portion of it. For example, if the *MPC* is 0.7, then households spend 70 cents of each additional dollar of disposable income on consumer goods and services and save 30 cents.


Figure 3-5 illustrates the consumption function. The slope of the consumption function tells us how much consumption increases when disposable income increases by one dollar. That is, the slope of the consumption function is the *MPC*.

Investment

Both firms and households purchase investment goods. Firms buy investment goods to add to their stock of capital and to replace existing capital as it wears out. Households buy new houses, which are also part of investment. Total investment in the United States averages about 15 percent of GDP.

The quantity of investment goods demanded depends on the interest rate, which measures the cost of the funds used to finance investment. For an investment project to be profitable, its return (the revenue from increased future

figure 3-5

**The Consumption Function**

The consumption function relates consumption C to disposable income $Y - T$. The marginal propensity to consume MPC is the amount by which consumption increases when disposable income increases by one dollar.

production of goods and services) must exceed its cost (the payments for borrowed funds). If the interest rate rises, fewer investment projects are profitable, and the quantity of investment goods demanded falls.


For example, suppose a firm is considering whether it should build a \$1 million factory that would yield a return of \$100,000 per year, or 10 percent. The firm compares this return to the cost of borrowing the \$1 million. If the interest rate is below 10 percent, the firm borrows the money in financial markets and makes the investment. If the interest rate is above 10 percent, the firm forgoes the investment opportunity and does not build the factory.

The firm makes the same investment decision even if it does not have to borrow the \$1 million but rather uses its own funds. The firm can always deposit this money in a bank or a money market fund and earn interest on it. Building the factory is more profitable than the deposit if and only if the interest rate is less than the 10 percent return on the factory.

A person wanting to buy a new house faces a similar decision. The higher the interest rate, the greater the cost of carrying a mortgage. A \$100,000 mortgage costs \$8,000 per year if the interest rate is 8 percent and \$10,000 per year if the interest rate is 10 percent. As the interest rate rises, the cost of owning a home rises, and the demand for new homes falls.

When studying the role of interest rates in the economy, economists distinguish between the nominal interest rate and the real interest rate. This distinction is relevant when the overall level of prices is changing. The **nominal interest rate** is the interest rate as usually reported: it is the rate of interest that investors pay to borrow money. The **real interest rate** is the nominal interest rate corrected for the effects of inflation. If the nominal interest rate is 8 percent and the inflation rate is 3 percent, then the real interest rate is 5 percent. In Chapter 4 we discuss the relation between nominal and real interest rates in detail. Here it is sufficient to note that the real interest rate measures the true cost of borrowing and, thus, determines the quantity of investment.

figure 3-6

**The Investment Function**

The investment function relates the quantity of investment I to the real interest rate r . Investment depends on the real interest rate because the interest rate is the cost of borrowing. The investment function slopes downward: when the interest rate rises, fewer investment projects are profitable.

We can summarize this discussion with an equation relating investment I to the real interest rate r :

$$I = I(r).$$

Figure 3-6 shows this investment function. It slopes downward, because as the interest rate rises, the quantity of investment demanded falls.

Government Purchases

Government purchases are the third component of the demand for goods and services. The federal government buys guns, missiles, and the services of government employees. Local governments buy library books, build schools, and hire teachers. Governments at all levels build roads and other public works. All these transactions make up government purchases of goods and services, which account for about 20 percent of GDP in the United States.

These purchases are only one type of government spending. The other type is transfer payments to households, such as welfare for the poor and Social Security payments for the elderly. Unlike government purchases, transfer payments are not made in exchange for some of the economy's output of goods and services. Therefore, they are not included in the variable G .

Transfer payments do affect the demand for goods and services indirectly. Transfer payments are the opposite of taxes: they increase households' disposable income, just as taxes reduce disposable income. Thus, an increase in transfer payments financed by an increase in taxes leaves disposable income unchanged. We can now revise our definition of T to equal taxes minus transfer payments. Disposable income, $Y - T$, includes both the negative impact of taxes and the positive impact of transfer payments.

If government purchases equal taxes minus transfers, then $G = T$, and the government has a *balanced budget*. If G exceeds T , the government runs a *budget*



The Many Different Interest Rates

If you look in the business section of a newspaper, you will find many different interest rates reported. By contrast, throughout this book, we will talk about “the” interest rate, as if there were only one interest rate in the economy. The only distinction we will make is between the nominal interest rate (which is not corrected for inflation) and the real interest rate (which is corrected for inflation). Almost all of the interest rates reported in the newspaper are nominal.

Why does the newspaper report so many interest rates? The various interest rates differ in three ways:

- **Term.** Some loans in the economy are for short periods of time, even as short as overnight. Other loans are for 30 years or even longer. The interest rate on a loan depends on its term. Long-term interest rates are usually, but not always, higher than short-term interest rates.
- **Credit risk.** In deciding whether to make a loan, a lender must take into account the probability that the borrower will repay. The law allows borrowers to default on their loans by declaring bankruptcy. The higher the perceived prob-

ability of default, the higher the interest rate. The safest credit risk is the government, and so government bonds tend to pay a low interest rate. At the other extreme, financially shaky corporations can raise funds only by issuing *junk bonds*, which pay a high interest rate to compensate for the high risk of default.

- **Tax treatment.** The interest on different types of bonds is taxed differently. Most important, when state and local governments issue bonds, called *municipal bonds*, the holders of the bonds do not pay federal income tax on the interest income. Because of this tax advantage, municipal bonds pay a lower interest rate.

When you see two different interest rates in the newspaper, you can almost always explain the difference by considering the term, the credit risk, and the tax treatment of the loan.

Although there are many different interest rates in the economy, macroeconomists can usually ignore these distinctions. The various interest rates tend to move up and down together. The assumption that there is only one interest rate is, for our purposes, a useful simplification.

deficit, which it funds by issuing government debt—that is, by borrowing in the financial markets. If G is less than T , the government runs a *budget surplus*, which it can use to repay some of its outstanding debt.

Here we do not try to explain the political process that leads to a particular fiscal policy—that is, to the level of government purchases and taxes. Instead, we take government purchases and taxes as exogenous variables. To denote that these variables are fixed outside of our model of national income, we write

$$G = \bar{G}.$$

$$T = \bar{T}.$$

We do, however, want to examine the impact of fiscal policy on the variables determined within the model, the endogenous variables. The endogenous variables here are consumption, investment, and the interest rate.

To see how the exogenous variables affect the endogenous variables, we must complete the model. This is the subject of the next section.

3-4 What Brings the Supply and Demand for Goods and Services Into Equilibrium?

We have now come full circle in the circular flow diagram, Figure 3-1. We began by examining the supply of goods and services, and we have just discussed the demand for them. How can we be certain that all these flows balance? In other words, what ensures that the sum of consumption, investment, and government purchases equals the amount of output produced? We will see that in this classical model, the interest rate has the crucial role of equilibrating supply and demand.

There are two ways to think about the role of the interest rate in the economy. We can consider how the interest rate affects the supply and demand for goods or services. Or we can consider how the interest rate affects the supply and demand for loanable funds. As we will see, these two approaches are two sides of the same coin.

Equilibrium in the Market for Goods and Services: The Supply and Demand for the Economy's Output

The following equations summarize the discussion of the demand for goods and services in Section 3-3:

$$Y = C + I + G.$$

$$C = C(Y - T).$$

$$I = I(r).$$

$$G = \bar{G}.$$

$$T = \bar{T}.$$

The demand for the economy's output comes from consumption, investment, and government purchases. Consumption depends on disposable income; investment depends on the real interest rate; and government purchases and taxes are the exogenous variables set by fiscal policymakers.

To this analysis, let's add what we learned about the supply of goods and services in Section 3-1. There we saw that the factors of production and the production function determine the quantity of output supplied to the economy:

$$\begin{aligned} Y &= F(\bar{K}, \bar{L}) \\ &= \bar{Y}. \end{aligned}$$

Now let's combine these equations describing the supply and demand for output. If we substitute the consumption function and the investment function into the national income accounts identity, we obtain

$$Y = C(Y - T) + I(r) + G.$$

Because the variables G and T are fixed by policy, and the level of output Y is fixed by the factors of production and the production function, we can write

$$\bar{Y} = C(\bar{Y} - \bar{T}) + I(r) + \bar{G}.$$

This equation states that the supply of output equals its demand, which is the sum of consumption, investment, and government purchases.

Notice that the interest rate r is the only variable not already determined in the last equation. This is because the interest rate still has a key role to play: it must adjust to ensure that the demand for goods equals the supply. The greater the interest rate, the lower the level of investment, and thus the lower the demand for goods and services, $C + I + G$. If the interest rate is too high, investment is too low, and the demand for output falls short of the supply. If the interest rate is too low, investment is too high, and the demand exceeds the supply. *At the equilibrium interest rate, the demand for goods and services equals the supply.*

This conclusion may seem somewhat mysterious. One might wonder how the interest rate gets to the level that balances the supply and demand for goods and services. The best way to answer this question is to consider how financial markets fit into the story.

Equilibrium in the Financial Markets: The Supply and Demand for Loanable Funds

Because the interest rate is the cost of borrowing and the return to lending in financial markets, we can better understand the role of the interest rate in the economy by thinking about the financial markets. To do this, rewrite the national income accounts identity as

$$Y - C - G = I.$$

The term $Y - C - G$ is the output that remains after the demands of consumers and the government have been satisfied; it is called **national saving** or simply **saving** (S). In this form, the national income accounts identity shows that saving equals investment.

To understand this identity more fully, we can split national saving into two parts—one part representing the saving of the private sector and the other representing the saving of the government:

$$(Y - T - C) + (T - G) = I.$$

The term $(Y - T - C)$ is disposable income minus consumption, which is **private saving**. The term $(T - G)$ is government revenue minus government spending, which is **public saving**. (If government spending exceeds government revenue, the government runs a budget deficit, and public saving is negative.) National saving is the sum of private and public saving. The circular flow

diagram in Figure 3-1 reveals an interpretation of this equation: this equation states that the flows into the financial markets (private and public saving) must balance the flows out of the financial markets (investment).

To see how the interest rate brings financial markets into equilibrium, substitute the consumption function and the investment function into the national income accounts identity:

$$Y - C(Y - T) - G = I(r).$$

Next, note that G and T are fixed by policy and Y is fixed by the factors of production and the production function:


$$\begin{aligned}\bar{Y} - C(\bar{Y} - \bar{T}) - \bar{G} &= I(r) \\ \bar{S} &= I(r).\end{aligned}$$

The left-hand side of this equation shows that national saving depends on income Y and the fiscal-policy variables G and T . For fixed values of Y , G , and T , national saving S is also fixed. The right-hand side of the equation shows that investment depends on the interest rate.

Figure 3-7 graphs saving and investment as a function of the interest rate. The saving function is a vertical line because in this model saving does not depend on the interest rate (although we relax this assumption later). The investment function slopes downward: the higher the interest rate, the fewer profitable investment projects.

From a quick glance at Figure 3-7, one might think it was a supply-and-demand diagram for a particular good. In fact, saving and investment can be interpreted in terms of supply and demand. In this case, the “good” is **loanable funds**, and its “price” is the interest rate. Saving is the supply of loanable funds—

figure 3-7



Saving, Investment, and the Interest Rate The interest rate adjusts to bring saving and investment into balance. The vertical line represents saving—the supply of loanable funds. The downward-sloping line represents investment—the demand for loanable funds. The intersection of these two curves determines the equilibrium interest rate.

households lend their saving to investors or deposit their saving in a bank that then loans the funds out. Investment is the demand for loanable funds—investors borrow from the public directly by selling bonds or indirectly by borrowing from banks. Because investment depends on the interest rate, the quantity of loanable funds demanded also depends on the interest rate.

The interest rate adjusts until the amount that firms want to invest equals the amount that households want to save. If the interest rate is too low, investors want more of the economy's output than households want to save. Equivalently, the quantity of loanable funds demanded exceeds the quantity supplied. When this happens, the interest rate rises. Conversely, if the interest rate is too high, households want to save more than firms want to invest; because the quantity of loanable funds supplied is greater than the quantity demanded, the interest rate falls. The equilibrium interest rate is found where the two curves cross. *At the equilibrium interest rate, households' desire to save balances firms' desire to invest, and the quantity of loanable funds supplied equals the quantity demanded.*

Changes in Saving: The Effects of Fiscal Policy


We can use our model to show how fiscal policy affects the economy. When the government changes its spending or the level of taxes, it affects the demand for the economy's output of goods and services and alters national saving, investment, and the equilibrium interest rate.

An Increase in Government Purchases Consider first the effects of an increase in government purchases of an amount ΔG . The immediate impact is to increase the demand for goods and services by ΔG . But since total output is fixed by the factors of production, the increase in government purchases must be met by a decrease in some other category of demand. Because disposable income $Y - T$ is unchanged, consumption C is unchanged. The increase in government purchases must be met by an equal decrease in investment.

To induce investment to fall, the interest rate must rise. Hence, the increase in government purchases causes the interest rate to increase and investment to decrease. Government purchases are said to **crowd out** investment.

To grasp the effects of an increase in government purchases, consider the impact on the market for loanable funds. Because the increase in government purchases is not accompanied by an increase in taxes, the government finances the additional spending by borrowing—that is, by reducing public saving. With private saving unchanged, this government borrowing reduces national saving. As Figure 3-8 shows, a reduction in national saving is represented by a leftward shift in the supply of loanable funds available for investment. At the initial interest rate, the demand for loanable funds exceeds the supply. The equilibrium interest rate rises to the point where the investment schedule crosses the new saving schedule. Thus, an increase in government purchases causes the interest rate to rise from r_1 to r_2 .

figure 3-8



A Reduction in Saving A reduction in saving, possibly the result of a change in fiscal policy, shifts the saving schedule to the left. The new equilibrium is the point at which the new saving schedule crosses the investment schedule. A reduction in saving lowers the amount of investment and raises the interest rate. Fiscal-policy actions that reduce saving are said to crowd out investment.

CASE STUDY

Wars and Interest Rates in the United Kingdom, 1730–1920


Wars are traumatic—both for those who fight them and for a nation’s economy. Because the economic changes accompanying them are often large, wars provide a natural experiment with which economists can test their theories. We can learn about the economy by seeing how in wartime the endogenous variables respond to the major changes in the exogenous variables.

One exogenous variable that changes substantially in wartime is the level of government purchases. Figure 3-9 shows military spending as a percentage of GDP for the United Kingdom from 1730 to 1919. This graph shows, as one would expect, that government purchases rose suddenly and dramatically during the eight wars of this period.

Our model predicts that this wartime increase in government purchases—and the increase in government borrowing to finance the wars—should have raised the demand for goods and services, reduced the supply of loanable funds, and raised the interest rate. To test this prediction, Figure 3-9 also shows the interest rate on long-term government bonds, called *consols* in the United Kingdom. A positive association between military purchases and interest rates is apparent in this figure. These data support the model’s prediction: interest rates do tend to rise when government purchases increase.⁴

⁴ Daniel K. Benjamin and Levis A. Kochin, “War, Prices, and Interest Rates: A Martial Solution to Gibson’s Paradox,” in M. D. Bordo and A. J. Schwartz, eds., *A Retrospective on the Classical Gold Standard, 1821–1931* (Chicago: University of Chicago Press, 1984), 587–612; Robert J. Barro, “Government Spending, Interest Rates, Prices, and Budget Deficits in the United Kingdom, 1701–1918,” *Journal of Monetary Economics* 20 (September 1987): 221–248.

figure 3-9



Military Spending and the Interest Rate in the United Kingdom This figure shows military spending as a percentage of GDP in the United Kingdom from 1730 to 1919. Not surprisingly, military spending rose substantially during each of the eight wars of this period. This figure also shows that the interest rate tended to rise when military spending rose.

Source: Series constructed from various sources described in Robert J. Barro, "Government Spending, Interest Rates, Prices, and Budget Deficits in the United Kingdom, 1701–1918," *Journal of Monetary Economics* 20 (September 1987): 221–248.

One problem with using wars to test theories is that many economic changes may be occurring at the same time. For example, in World War II, while government purchases increased dramatically, rationing also restricted consumption of many goods. In addition, the risk of defeat in the war and default by the government on its debt presumably increases the interest rate the government must pay. Economic models predict what happens when one exogenous variable changes and all the other exogenous variables remain constant. In the real world, however, many exogenous variables may change at once. Unlike controlled laboratory experiments, the natural experiments on which economists must rely are not always easy to interpret.

A Decrease in Taxes Now consider a reduction in taxes of ΔT . The immediate impact of the tax cut is to raise disposable income and thus to raise consumption. Disposable income rises by ΔT , and consumption rises by an amount equal to

ΔT times the marginal propensity to consume MPC . The higher the MPC , the greater the impact of the tax cut on consumption.

Because the economy's output is fixed by the factors of production and the level of government purchases is fixed by the government, the increase in consumption must be met by a decrease in investment. For investment to fall, the interest rate must rise. Hence, a reduction in taxes, like an increase in government purchases, crowds out investment and raises the interest rate.

We can also analyze the effect of a tax cut by looking at saving and investment. Because the tax cut raises disposable income by ΔT , consumption goes up by $MPC \times \Delta T$. National saving S , which equals $Y - C - G$, falls by the same amount as consumption rises. As in Figure 3-8, the reduction in saving shifts the supply of loanable funds to the left, which increases the equilibrium interest rate and crowds out investment.

Changes in Investment Demand

So far, we have discussed how fiscal policy can change national saving. We can also use our model to examine the other side of the market—the demand for investment. In this section we look at the causes and effects of changes in investment demand.

One reason investment demand might increase is technological innovation. Suppose, for example, that someone invents a new technology, such as the railroad or the computer. Before a firm or household can take advantage of the innovation, it must buy investment goods. The invention of the railroad had no value until railroad cars were produced and tracks were laid. The idea of the computer was not productive until computers were manufactured. Thus, technological innovation leads to an increase in investment demand.


Investment demand may also change because the government encourages or discourages investment through the tax laws. For example, suppose that the government increases personal income taxes and uses the extra revenue to provide tax cuts for those who invest in new capital. Such a change in the tax laws makes more investment projects profitable and, like a technological innovation, increases the demand for investment goods.

Figure 3-10 shows the effects of an increase in investment demand. At any given interest rate, the demand for investment goods (and also for loanable funds) is higher. This increase in demand is represented by a shift in the investment schedule to the right. The economy moves from the old equilibrium, point A, to the new equilibrium, point B.

The surprising implication of Figure 3-10 is that the equilibrium amount of investment is unchanged. Under our assumptions, the fixed level of saving determines the amount of investment; in other words, there is a fixed supply of loanable funds. An increase in investment demand merely raises the equilibrium interest rate.

We would reach a different conclusion, however, if we modified our simple consumption function and allowed consumption (and its flip side, saving) to

figure 3-10




An Increase in the Demand for Investment

An increase in the demand for investment goods shifts the investment schedule to the right. At any given interest rate, the amount of investment is greater. The equilibrium moves from point A to point B. Because the amount of saving is fixed, the increase in investment demand raises the interest rate while leaving the equilibrium amount of investment unchanged.

depend on the interest rate. Because the interest rate is the return to saving (as well as the cost of borrowing), a higher interest rate might reduce consumption and increase saving. If so, the saving schedule would be upward sloping, rather than vertical.

With an upward-sloping saving schedule, an increase in investment demand would raise both the equilibrium interest rate and the equilibrium quantity of investment. Figure 3-11 shows such a change. The increase in the interest rate causes households to consume less and save more. The decrease in consumption frees resources for investment.

figure 3-11



An Increase in Investment Demand When Saving Depends on the Interest Rate

When saving is positively related to the interest rate, a rightward shift in the investment schedule increases the interest rate and the amount of investment. The higher interest rate induces people to increase saving, which in turn allows investment to increase.



The Identification Problem

In our model, investment depends on the interest rate. The higher the interest rate, the fewer investment projects there are that are profitable. The investment schedule therefore slopes downward.

Economists who look at macroeconomic data, however, usually fail to find an obvious association between investment and interest rates. In years when interest rates are high, investment is not always low. In years when interest rates are low, investment is not always high.

How do we interpret this finding? Does it mean that investment does not depend on the interest rate? Does it suggest that our model of saving, investment, and the interest rate is inconsistent with how the economy actually functions?

Luckily, we do not have to discard our model. The inability to find an empirical relationship between investment and interest rates is an example of the *identification problem*. The identification problem arises when variables are related in more than one way. When we look at data, we are observing a combination of these different relationships, and it is difficult to “identify” any one of them.

To understand this problem more concretely, consider the relationships among saving, investment, and the interest rate. Suppose, on the one hand, that all changes in the interest rate resulted from changes in saving—that is, from

shifts in the saving schedule. Then, as shown in the left-hand side of panel (a) in Figure 3-12, all changes would represent movement along a fixed investment schedule. As the right-hand side of panel (a) shows, the data would trace out this investment schedule. Thus, we would observe a negative relationship between investment and interest rates.


Suppose, on the other hand, that all changes in the interest rate resulted from technological innovations—that is, from shifts in the investment schedule. Then, as shown in panel (b), all changes would represent movements in the investment schedule along a fixed saving schedule. As the right-hand side of panel (b) shows, the data would reflect this saving schedule. Thus, we would observe a positive relationship between investment and interest rates.

In the real world, interest rates change sometimes because of shifts in the saving schedule and sometimes because of shifts in the investment schedule. In this mixed case, as shown in panel (c), a plot of the data would reveal no recognizable relation between interest rates and the quantity of investment, just as economists observe in actual data. The moral of the story is simple and is applicable to many other situations: the empirical relationship we expect to observe depends crucially on which exogenous variables we think are changing.

3-5 Conclusion

In this chapter we have developed a model that explains the production, distribution, and allocation of the economy’s output of goods and services. Because the model incorporates all the interactions illustrated in the circular flow diagram in Figure 3-1, it is sometimes called a *general equilibrium model*. The model emphasizes how prices adjust to equilibrate supply and demand. Factor prices equilibrate factor markets. The interest rate equilibrates the supply and demand for goods and services (or, equivalently, the supply and demand for loanable funds).

figure 3-12

**Identifying the Investment Function**

When we look at data on interest rates r and investment I , what we find depends on which exogenous variables are changing. In panel (a), the saving schedule is shifting, perhaps because of changes in fiscal policy; we would observe a negative correlation between r and I . In panel (b), the investment schedule is shifting, perhaps because of technological innovations; we would observe a positive correlation between r and I . In the more realistic situation shown in panel (c), both schedules are shifting. In the data, we would observe no correlation between r and I , which is in fact what researchers typically find.

Throughout the chapter, we have discussed various applications of the model. The model can explain how income is divided among the factors of production and how factor prices depend on factor supplies. We have also used the model to discuss how fiscal policy alters the allocation of output among its alternative uses—consumption, investment, and government purchases—and how it affects the equilibrium interest rate.

At this point it is useful to review some of the simplifying assumptions we have made in this chapter. In the following chapters we relax some of these assumptions in order to address a greater range of questions.

- We have ignored the role of money, the asset with which goods and services are bought and sold. In Chapter 4 we discuss how money affects the economy and the influence of monetary policy.
- We have assumed that there is no trade with other countries. In Chapter 5 we consider how international interactions affect our conclusions.
- We have assumed that the labor force is fully employed. In Chapter 6 we examine the reasons for unemployment and see how public policy influences the level of unemployment.
- We have assumed that the capital stock, the labor force, and the production technology are fixed. In Chapters 7 and 8 we see how changes over time in each of these lead to growth in the economy's output of goods and services.
- We have ignored the role of short-run sticky prices. In Chapters 9 through 13, we develop a model of short-run fluctuations that includes sticky prices. We then discuss how the model of short-run fluctuations relates to the model of national income developed in this chapter.

Before going on to these chapters, go back to the beginning of this one and make sure you can answer the four groups of questions about national income that begin the chapter.

Summary

1. The factors of production and the production technology determine the economy's output of goods and services. An increase in one of the factors of production or a technological advance raises output.
2. Competitive, profit-maximizing firms hire labor until the marginal product of labor equals the real wage. Similarly, these firms rent capital until the marginal product of capital equals the real rental price. Therefore, each factor of production is paid its marginal product. If the production function has constant returns to scale, all output is used to compensate the inputs.
3. The economy's output is used for consumption, investment, and government purchases. Consumption depends positively on disposable income. Investment depends negatively on the real interest rate. Government purchases and taxes are the exogenous variables of fiscal policy.
4. The real interest rate adjusts to equilibrate the supply and demand for the economy's output—or, equivalently, to equilibrate the supply of loanable funds (saving) and the demand for loanable funds (investment). A decrease in national saving, perhaps because of an increase in government purchases or a

decrease in taxes, reduces the equilibrium amount of investment and raises the interest rate. An increase in investment demand, perhaps because of a technological innovation or a tax incentive for investment, also raises the interest rate. An increase in investment demand increases the quantity of investment only if higher interest rates stimulate additional saving.

KEY CONCEPTS

Factors of production	Marginal product of capital (MPK)	Nominal interest rate
Production function	Real rental price of capital	Real interest rate
Constant returns to scale	Economic profit versus accounting profit	National saving (saving)
Factor prices	Disposable income	Private saving
Competition	Consumption function	Public saving
Marginal product of labor (MPL)	Marginal propensity to consume (MPC)	Loanable funds
Diminishing marginal product		Crowding out
Real wage		

QUESTIONS FOR REVIEW

1. What determines the amount of output an economy produces?
2. Explain how a competitive, profit-maximizing firm decides how much of each factor of production to demand.
3. What is the role of constant returns to scale in the distribution of income?
4. What determines consumption and investment?
5. Explain the difference between government purchases and transfer payments. Give two examples of each.
6. What makes the demand for the economy's output of goods and services equal the supply?
7. Explain what happens to consumption, investment, and the interest rate when the government increases taxes.

PROBLEMS AND APPLICATIONS

1. Use the neoclassical theory of distribution to predict the impact on the real wage and the real rental price of capital of each of the following events:
 - a. A wave of immigration increases the labor force.
 - b. An earthquake destroys some of the capital stock.
 - c. A technological advance improves the production function.
2. If a 10-percent increase in both capital and labor causes output to increase by less than 10 percent, the production function is said to exhibit *decreasing returns to scale*. If it causes output to increase by more than 10 percent, the production function is said to exhibit *increasing returns to scale*. Why might a production function exhibit decreasing or increasing returns to scale?
3. According to the neoclassical theory of distribution, the real wage earned by any worker equals that worker's marginal productivity. Let's use this insight to examine the incomes of two groups of workers: farmers and barbers.

- a. Over the past century, the productivity of farmers has risen substantially because of technological progress. According to the neoclassical theory, what should have happened to their real wage?
 - b. In what units is the real wage discussed in part (a) measured?
 - c. Over the same period, the productivity of barbers has remained constant. What should have happened to their real wage?
 - d. In what units is the real wage in part (c) measured?
 - e. Suppose workers can move freely between being farmers and being barbers. What does this mobility imply for the wages of farmers and barbers?
 - f. What do your previous answers imply for the price of haircuts relative to the price of food?
 - g. Who benefits from technological progress in farming—farmers or barbers?
4. The government raises taxes by \$100 billion. If the marginal propensity to consume is 0.6, what happens to the following? Do they rise or fall? By what amounts?
 - a. Public saving.
 - b. Private saving.
 - c. National saving.
 - d. Investment.
 5. Suppose that an increase in consumer confidence raises consumers' expectations of future income and thus the amount they want to consume today. This might be interpreted as an upward shift in the consumption function. How does this shift affect investment and the interest rate?
 6. Consider an economy described by the following equations:

$$Y = C + I + G,$$

$$Y = 5,000,$$

$$G = 1,000,$$

$$T = 1,000,$$

$$C = 250 + 0.75(Y - T),$$

$$I = 1,000 - 50r.$$

- a. In this economy, compute private saving, public saving, and national saving.
 - b. Find the equilibrium interest rate.
 - c. Now suppose that G rises to 1,250. Compute private saving, public saving, and national saving.
 - d. Find the new equilibrium interest rate.
7. Suppose that the government increases taxes and government purchases by equal amounts. What happens to the interest rate and investment in response to this balanced-budget change? Does your answer depend on the marginal propensity to consume?
 8. When the government subsidizes investment, such as with an investment tax credit, the subsidy often applies to only some types of investment. This question asks you to consider the effect of such a change. Suppose there are two types of investment in the economy: business investment and residential investment. And suppose that the government institutes an investment tax credit only for business investment.
 - a. How does this policy affect the demand curve for business investment? The demand curve for residential investment?
 - b. Draw the economy's supply and demand for loanable funds. How does this policy affect the supply and demand for loanable funds? What happens to the equilibrium interest rate?
 - c. Compare the old and the new equilibrium. How does this policy affect the total quantity of investment? The quantity of business investment? The quantity of residential investment?
 9. If consumption depended on the interest rate, how would that affect the conclusions reached in this chapter about the effects of fiscal policy?

The Cobb–Douglas Production Function

What production function describes how actual economies turn capital and labor into GDP? The answer to this question came from a historic collaboration between a U.S. senator and a mathematician.

Paul Douglas was a U.S. senator from Illinois from 1949 to 1966. In 1927, however, when he was still a professor of economics, he noticed a surprising fact: the division of national income between capital and labor had been roughly constant over a long period. In other words, as the economy grew more prosperous over time, the total income of workers and the total income of capital owners grew at almost exactly the same rate. This observation caused Douglas to wonder what conditions lead to constant factor shares.

Douglas asked Charles Cobb, a mathematician, what production function, if any, would produce constant factor shares if factors always earned their marginal products. The production function would need to have the property that

$$\text{Capital Income} = MPK \times K = \alpha Y$$

and

$$\text{Labor Income} = MPL \times L = (1 - \alpha) Y,$$

where α is a constant between zero and one that measures capital's share of income. That is, α determines what share of income goes to capital and what share goes to labor. Cobb showed that the function with this property is

$$Y = F(K, L) = AK^\alpha L^{1-\alpha},$$

where A is a parameter greater than zero that measures the productivity of the available technology. This function became known as the *Cobb–Douglas production function*.

Let's take a closer look at some of the properties of this production function. First, the Cobb–Douglas production function has constant returns to scale. That is, if capital and labor are increased by the same proportion, then output increases by that proportion as well.⁵

⁵ *Mathematical note:* To prove that the Cobb–Douglas production function has constant returns to scale, examine what happens when we multiply capital and labor by a constant z :

$$F(zK, zL) = A(zK)^\alpha (zL)^{1-\alpha}.$$

Expanding terms on the right,

$$F(zK, zL) = Az^\alpha K^\alpha z^{1-\alpha} L^{1-\alpha}.$$

(footnote continues)

Next, consider the marginal products for the Cobb–Douglas production function. The marginal product of labor is⁶

$$MPL = (1 - \alpha) AK^\alpha L^{-\alpha},$$

and the marginal product of capital is

$$MPK = \alpha AK^{\alpha-1} L^{1-\alpha}.$$

From these equations, recalling that α is between zero and one, we can see what causes the marginal products of the two factors to change. An increase in the amount of capital raises the MPL and reduces the MPK . Similarly, an increase in the amount of labor reduces the MPL and raises the MPK . A technological advance that increases the parameter A raises the marginal product of both factors proportionately.

The marginal products for the Cobb–Douglas production function can also be written as⁷

$$MPL = (1 - \alpha) Y/L.$$

$$MPK = \alpha Y/K.$$

The MPL is proportional to output per worker, and the MPK is proportional to output per unit of capital. Y/L is called *average labor productivity*, and Y/K is called *average capital productivity*. If the production function is Cobb–Douglas, then the marginal productivity of a factor is proportional to its average productivity.

We can now verify that if factors earn their marginal products, then the parameter α indeed tells us how much income goes to labor and how much goes to capital. The total wage bill, which we have seen is $MPL \times L$, is simply $(1 - \alpha)Y$.

Rearranging to bring like terms together, we get

$$F(zK, zL) = z^\alpha z^{1-\alpha} AK^\alpha L^{1-\alpha}.$$

Since $z^\alpha z^{1-\alpha} = z$, our function becomes

$$F(zK, zL) = zAK^\alpha L^{1-\alpha}.$$

But $AK^\alpha L^{1-\alpha} = F(K, L)$. Thus,

$$F(zK, zL) = zF(K, L) = zY.$$

Hence, the amount of output Y increases by the same factor z , which implies that this production function has constant returns to scale.


⁶ *Mathematical note:* Obtaining the formulas for the marginal products from the production function requires a bit of calculus. To find the MPL , differentiate the production function with respect to L . This is done by multiplying by the exponent $(1 - \alpha)$, and then subtracting 1 from the old exponent to obtain the new exponent, $-\alpha$. Similarly, to obtain the MPK , differentiate the production function with respect to K .

⁷ *Mathematical note:* To check these expressions for the marginal products, substitute in the production function for Y to show that these expressions are equivalent to the earlier formulas for the marginal products.

Therefore, $(1 - \alpha)$ is labor's share of output. Similarly, the total return to capital, $MPK \times K$, is αY , and α is capital's share of output. The ratio of labor income to capital income is a constant, $(1 - \alpha)/\alpha$, just as Douglas observed. The factor shares depend only on the parameter α , not on the amounts of capital or labor or on the state of technology as measured by the parameter A .

More recent U.S. data are also consistent with the Cobb–Douglas production function. Figure 3-13 shows the ratio of labor income to total income in the United States from 1960 to 2000. Despite the many changes in the economy over the past four decades, this ratio has remained about 0.7. This division of income is easily explained by a Cobb–Douglas production function in which the parameter α is about 0.3.

figure 3-13



The Ratio of Labor Income to Total Income Labor income has remained about 0.7 of total income over a long period of time. This approximate constancy of factor shares is evidence for the Cobb–Douglas production function. (This figure is produced from U.S. national income accounts data. Labor income is compensation of employees. Total income is the sum of labor income, corporate profits, net interest, rental income, and depreciation. Proprietors' income is excluded from these calculations, because it is a combination of labor income and capital income.)

Source: U.S. Department of Commerce.

MORE PROBLEMS AND APPLICATIONS

1. Suppose that the production function is Cobb–Douglas with parameter $\alpha = 0.3$.
 - a. What fractions of income do capital and labor receive?
 - b. Suppose that immigration raises the labor force by 10 percent. What happens to total output (in percent)? The rental price of capital? The real wage?
 - c. Suppose that a gift of capital from abroad raises the capital stock by 10 percent. What happens to total output (in percent)? The rental price of capital? The real wage?
 - d. Suppose that a technological advance raises the value of the parameter A by 10 percent. What happens to total output (in percent)? The rental price of capital? The real wage?
2. (This problem requires the use of calculus.) Consider a Cobb–Douglas production function with three inputs. K is capital (the number of machines), L is labor (the number of workers), and H is human capital (the number of college degrees among the workers). The production function is

$$Y = K^{1/3} L^{1/3} H^{1/3}.$$

- a. Derive an expression for the marginal product of labor. How does an increase in the amount

of human capital affect the marginal product of labor?

- b. Derive an expression for the marginal product of human capital. How does an increase in the amount of human capital affect the marginal product of human capital?
- c. What is the income share paid to labor? What is the income share paid to human capital? In the national income accounts of this economy, what share of total income do you think workers would appear to receive? (*Hint:* Consider where the return to human capital shows up.)
- d. An unskilled worker earns the marginal product of labor, whereas a skilled worker earns the marginal product of labor plus the marginal product of human capital. Using your answers to (a) and (b), find the ratio of the skilled wage to the unskilled wage. How does an increase in the amount of human capital affect this ratio? Explain.
- e. Some people advocate government funding of college scholarships as a way of creating a more egalitarian society. Others argue that scholarships help only those who are able to go to college. Do your answers to the preceding questions shed light on this debate?

Money and Inflation

There is no subtler, no surer means of overturning the existing basis of society than to debauch the currency. The process engages all the hidden forces of economic law on the side of destruction, and does it in a manner which not one man in a million is able to diagnose.

— John Maynard Keynes

In 1970 the *New York Times* cost 15 cents, the median price of a single-family home was \$23,400, and the average wage in manufacturing was \$3.36 per hour. In 2000 the *Times* cost 75 cents, the price of a home was \$166,000, and the average wage was \$14.26 per hour. This overall increase in prices is called **inflation**, and it is the subject of this chapter.

The rate of inflation—the percentage change in the overall level of prices—varies greatly over time and across countries. In the United States, according to the consumer price index, prices rose an average of 2.4 percent per year in the 1960s, 7.1 percent per year in the 1970s, 5.5 percent per year in the 1980s, and 3.0 percent in the 1990s. Even when the U.S. inflation problem became severe during the 1970s, it was nothing compared to the episodes of extraordinarily high inflation, called **hyperinflation**, that other countries have experienced from time to time. A classic example is Germany in 1923, when prices rose an average of 500 percent *per month*.

In this chapter we examine the classical theory of the causes, effects, and social costs of inflation. The theory is “classical” in the sense that it assumes that prices are flexible. As we first discussed in Chapter 1, most economists believe this assumption describes the behavior of the economy in the long run. By contrast, many prices are thought to be sticky in the short run, and beginning in Chapter 9, we incorporate this fact into our analysis. Yet, for now, we ignore short-run price stickiness. As we will see, the classical theory of inflation not only provides a good description of the long run, it also provides a useful foundation for the short-run analysis we develop later.

The “hidden forces of economic law” that lead to inflation are not as mysterious as Keynes claims in the quotation that opens this chapter. Inflation is simply an increase in the average level of prices, and a price is the rate at which money is exchanged for a good or a service. To understand inflation, therefore,

we must understand money—what it is, what affects its supply and demand, and what influence it has on the economy. Thus, Section 4-1 begins our analysis of inflation by discussing the economist’s concept of “money” and how, in most modern economies, the government controls the quantity of money in the hands of the public. Section 4-2 shows that the quantity of money determines the price level and that the rate of growth in the quantity of money determines the rate of inflation.

Inflation in turn has numerous effects of its own on the economy. Section 4-3 discusses the revenue that the government raises by printing money, sometimes called the *inflation tax*. Section 4-4 examines how inflation affects the nominal interest rate. Section 4-5 discusses how the nominal interest rate in turn affects the quantity of money people wish to hold and, thereby, the price level.

After completing our analysis of the causes and effects of inflation, in Section 4-6 we address what is perhaps the most important question about inflation: Is it a major social problem? Does inflation amount to “overturning the existing basis of society,” as the chapter’s opening quotation suggests?

Finally, in Section 4-7, we discuss the extreme case of hyperinflation. Hyperinflations are interesting to examine because they show clearly the causes, effects, and costs of inflation. Just as seismologists learn much by studying earthquakes, economists learn much by studying how hyperinflations begin and end.

4-1 What Is Money?

When we say that a person has a lot of money, we usually mean that he or she is wealthy. By contrast, economists use the term *money* in a more specialized way. To an economist, money does not refer to all wealth but only to one type of it: **money** is the stock of assets that can be readily used to make transactions. Roughly speaking, the dollars in the hands of the public make up the nation’s stock of money.

The Functions of Money

Money has three purposes. It is a store of value, a unit of account, and a medium of exchange.

As a **store of value**, money is a way to transfer purchasing power from the present to the future. If I work today and earn \$100, I can hold the money and spend it tomorrow, next week, or next month. Of course, money is an imperfect store of value: if prices are rising, the amount you can buy with any given quantity of money is falling. Even so, people hold money because they can trade the money for goods and services at some time in the future.

As a **unit of account**, money provides the terms in which prices are quoted and debts are recorded. Microeconomics teaches us that resources are allocated according to relative prices—the prices of goods relative to other goods—yet

stores post their prices in dollars and cents. A car dealer tells you that a car costs \$20,000, not 400 shirts (even though it may amount to the same thing). Similarly, most debts require the debtor to deliver a specified number of dollars in the future, not a specified amount of some commodity. Money is the yardstick with which we measure economic transactions.

As a **medium of exchange**, money is what we use to buy goods and services. “This note is legal tender for all debts, public and private” is printed on the U.S. dollar. When we walk into stores, we are confident that the shopkeepers will accept our money in exchange for the items they are selling. The ease with which money is converted into other things—goods and services—is sometimes called money’s *liquidity*.

To better understand the functions of money, try to imagine an economy without it: a barter economy. In such a world, trade requires the *double coincidence of wants*—the unlikely happenstance of two people each having a good that the other wants at the right time and place to make an exchange. A barter economy permits only simple transactions.

Money makes more indirect transactions possible. A professor uses her salary to buy books; the book publisher uses its revenue from the sale of books to buy paper; the paper company uses its revenue from the sale of paper to pay the lumberjack; the lumberjack uses his income to send his child to college; and the college uses its tuition receipts to pay the salary of the professor. In a complex, modern economy, trade is often indirect and requires the use of money.

The Types of Money

Money takes many forms. In the U.S. economy we make transactions with an item whose sole function is to act as money: dollar bills. These pieces of green paper with small portraits of famous Americans would have little value if they were not widely accepted as money. Money that has no intrinsic value is called **fiat money** because it is established as money by government decree, or fiat.

Although fiat money is the norm in most economies today, most societies in the past have used for money a commodity with some intrinsic value. Money of this sort is called **commodity money**. The most widespread example of commodity money is gold. When people use gold as money (or use paper money that is redeemable for gold), the economy is said to be on a **gold standard**. Gold is a form of commodity money because it can be used for various purposes—jewelry, dental fillings, and so on—as well as for transactions. The gold standard was common throughout the world during the late nineteenth century.



Drawing by Bernard Schoenbaum; © 1979
The New Yorker Magazine, Inc.

“And how would you like your funny money?”

CASE STUDY

Money in a POW Camp

An unusual form of commodity money developed in some Nazi prisoner of war (POW) camps during World War II. The Red Cross supplied the prisoners with various goods—food, clothing, cigarettes, and so on. Yet these rations were allocated without close attention to personal preferences, so the allocations were often inefficient. One prisoner may have preferred chocolate, while another may have preferred cheese, and a third may have wanted a new shirt. The differing tastes and endowments of the prisoners led them to trade with one another.

Barter proved to be an inconvenient way to allocate these resources, however, because it required the double coincidence of wants. In other words, a barter system was not the easiest way to ensure that each prisoner received the goods he valued most. Even the limited economy of the POW camp needed some form of money to facilitate transactions.

Eventually, cigarettes became the established “currency” in which prices were quoted and with which trades were made. A shirt, for example, cost about 80 cigarettes. Services were also quoted in cigarettes: some prisoners offered to do other prisoners’ laundry for 2 cigarettes per garment. Even nonsmokers were happy to accept cigarettes in exchange, knowing they could trade the cigarettes in the future for some good they did enjoy. Within the POW camp the cigarette became the store of value, the unit of account, and the medium of exchange.¹

How Fiat Money Evolves

It is not surprising that some form of commodity money arises to facilitate exchange: people are willing to accept a commodity currency such as gold because it has intrinsic value. The development of fiat money, however, is more perplexing. What would make people begin to value something that is intrinsically useless?

To understand how the evolution from commodity money to fiat money takes place, imagine an economy in which people carry around bags of gold. When a purchase is made, the buyer measures out the appropriate amount of gold. If the seller is convinced that the weight and purity of the gold are right, the buyer and seller make the exchange.

The government might first get involved in the monetary system to help people reduce transaction costs. Using raw gold as money is costly because it takes time to verify the purity of the gold and to measure the correct quantity. To reduce these costs, the government can mint gold coins of known purity and weight. The coins are easier to use than gold bullion because their values are widely recognized.

The next step is for the government to accept gold from the public in exchange for gold certificates—pieces of paper that can be redeemed for a certain

¹ R. A. Radford, “The Economic Organisation of a P.O.W. Camp,” *Economica* (November 1945): 189–201. The use of cigarettes as money is not limited to this example. In the Soviet Union in the late 1980s, packs of Marlboros were preferred to the ruble in the large underground economy.

quantity of gold. If people believe the government's promise to redeem the paper bills for gold, the bills are just as valuable as the gold itself. In addition, because the bills are lighter than gold (and gold coins), they are easier to use in transactions. Eventually, no one carries gold around at all, and these gold-backed government bills become the monetary standard.

Finally, the gold backing becomes irrelevant. If no one ever bothers to redeem the bills for gold, no one cares if the option is abandoned. As long as everyone continues to accept the paper bills in exchange, they will have value and serve as money. Thus, the system of commodity money evolves into a system of fiat money. Notice that in the end, the use of money in exchange is a social convention: everyone values fiat money because they expect everyone else to value it.

CASE STUDY

Money and Social Conventions on the Island of Yap

The economy of Yap, a small island in the Pacific, once had a type of money that was something between commodity and fiat money. The traditional medium of exchange in Yap was *fei*, stone wheels up to 12 feet in diameter. These stones had holes in the center so that they could be carried on poles and used for exchange.

Large stone wheels are not a convenient form of money. The stones were heavy, so it took substantial effort for a new owner to take his *fei* home after completing a transaction. Although the monetary system facilitated exchange, it did so at great cost.

Eventually, it became common practice for the new owner of the *fei* not to bother to take physical possession of the stone. Instead, the new owner accepted a claim to the *fei* without moving it. In future bargains, he traded this claim for goods that he wanted. Having physical possession of the stone became less important than having legal claim to it.

This practice was put to a test when a valuable stone was lost at sea during a storm. Because the owner lost his money by accident rather than through negligence, everyone agreed that his claim to the *fei* remained valid. Even generations later, when no one alive had ever seen this stone, the claim to this *fei* was still valued in exchange.²

How the Quantity of Money Is Controlled

The quantity of money available is called the **money supply**. In an economy that uses commodity money, the money supply is the quantity of that commodity. In an economy that uses fiat money, such as most economies today, the government controls the supply of money: legal restrictions give the government a monopoly on the printing of money. Just as the level of taxation and the level of government purchases are policy instruments of the government, so is the supply of money. The control over the money supply is called **monetary policy**.

² Norman Angell, *The Story of Money* (New York: Frederick A. Stokes Company, 1929), 88–89.

In the United States and many other countries, monetary policy is delegated to a partially independent institution called the **central bank**. The central bank of the United States is the **Federal Reserve**—often called *the Fed*. If you look at a U.S. dollar bill, you will see that it is called a *Federal Reserve Note*. Decisions over monetary policy are made by the Federal Open Market Committee. This committee is made up of members of the Federal Reserve Board, who are appointed by the president and confirmed by Congress, together with the presidents of the regional Federal Reserve Banks. The Federal Open Market Committee meets about every six weeks to discuss and set monetary policy.

The primary way in which the Fed controls the supply of money is through **open-market operations**—the purchase and sale of government bonds. When the Fed wants to increase the money supply, it uses some of the dollars it has to buy government bonds from the public. Because these dollars leave the Fed and enter into the hands of the public, the purchase increases the quantity of money in circulation. Conversely, when the Fed wants to decrease the money supply, it sells some government bonds from its own portfolio. This open-market sale of bonds takes some dollars out of the hands of the public and, thus, decreases the quantity of money in circulation.

In Chapter 18 we discuss in detail how the Fed controls the supply of money. For our current discussion, these details are not crucial. It is sufficient to assume that the Fed (or any other central bank) directly controls the supply of money.

How the Quantity of Money Is Measured

One goal of this chapter is to determine how the money supply affects the economy; we turn to that problem in the next section. As background for that analysis, let's first discuss how economists measure the quantity of money.

Because money is the stock of assets used for transactions, the quantity of money is the quantity of those assets. In simple economies, this quantity is easy to measure. In the POW camp, the quantity of money was the quantity of cigarettes in the camp. But how can we measure the quantity of money in more complex economies such as ours? The answer is not obvious, because no single asset is used for all transactions. People can use various assets, such as cash or checks, to make transactions, although some assets are more convenient than others. This ambiguity leads to numerous measures of the quantity of money.

The most obvious asset to include in the quantity of money is **currency**, the sum of outstanding paper money and coins. Most day-to-day transactions use currency as the medium of exchange.

A second type of asset used for transactions is **demand deposits**, the funds people hold in their checking accounts. If most sellers accept personal checks, assets in a checking account are almost as convenient as currency. In both cases, the assets are in a form ready to facilitate a transaction. Demand deposits are, therefore, added to currency when measuring the quantity of money.

Once we admit the logic of including demand deposits in the measured money stock, many other assets become candidates for inclusion. Funds in savings accounts, for example, can be easily transferred into checking accounts; these

assets are almost as convenient for transactions. Money market mutual funds allow investors to write checks against their accounts, although restrictions sometimes apply with regard to the size of the check or the number of checks written. Since these assets can be easily used for transactions, they should arguably be included in the quantity of money.

table 4-1**The Measures of Money**

Symbol	Assets Included	Amount in March 2001 (billions of dollars)
C	Currency	\$ 539
M1	Currency plus demand deposits, traveler's checks, and other checkable deposits	1,111
M2	M1 plus retail money market mutual fund balances, saving deposits (including money market deposit accounts), and small time deposits	5,100
M3	M2 plus large time deposits, repurchase agreements, Eurodollars, and institution-only money market mutual fund balances	7,326

Source: Federal Reserve.

Because it is hard to judge which assets should be included in the money stock, various measures are available. Table 4-1 presents the four measures of the money stock that the Federal Reserve calculates for the U.S. economy, together with a list of which assets are included in each measure. From the smallest to the largest, they are designated *C*, *M1*, *M2*, and *M3*. The most common measures for studying the effects of money on the economy are *M1* and *M2*. There is no consensus, however, about which measure of the money stock is best. Disagreements about monetary policy sometimes arise because different measures of money are moving in different directions.

4-2 The Quantity Theory of Money

Having defined what money is and described how it is controlled and measured, we can now examine how the quantity of money affects the economy. To do this, we must see how the quantity of money is related to other economic variables, such as prices and incomes.

Transactions and the Quantity Equation

People hold money to buy goods and services. The more money they need for such transactions, the more money they hold. Thus, the quantity of money in the economy is related to the number of dollars exchanged in transactions.

The link between transactions and money is expressed in the following equation, called the **quantity equation**:

$$\begin{array}{ccccccc} \text{Money} & \times & \text{Velocity} & = & \text{Price} & \times & \text{Transactions} \\ M & \times & V & = & P & \times & T. \end{array}$$

Let's examine each of the four variables in this equation.

The right-hand side of the quantity equation tells us about transactions. T represents the total number of transactions during some period of time, say, a year. In other words, T is the number of times in a year that goods or services are exchanged for money. P is the price of a typical transaction—the number of dollars exchanged. The product of the price of a transaction and the number of transactions, PT , equals the number of dollars exchanged in a year.

The left-hand side of the quantity equation tells us about the money used to make the transactions. M is the quantity of money. V is called the **transactions velocity of money** and measures the rate at which money circulates in the economy. In other words, velocity tells us the number of times a dollar bill changes hands in a given period of time.

For example, suppose that 60 loaves of bread are sold in a given year at \$0.50 per loaf. Then T equals 60 loaves per year, and P equals \$0.50 per loaf. The total number of dollars exchanged is

$$PT = \$0.50/\text{loaf} \times 60 \text{ loaves/year} = \$30/\text{year}.$$

The right-hand side of the quantity equation equals \$30 per year, which is the dollar value of all transactions.

Suppose further that the quantity of money in the economy is \$10. By rearranging the quantity equation, we can compute velocity as

$$\begin{aligned} V &= PT/M \\ &= (\$30/\text{year})/(\$10) \\ &= 3 \text{ times per year.} \end{aligned}$$

That is, for \$30 of transactions per year to take place with \$10 of money, each dollar must change hands 3 times per year.

The quantity equation is an *identity*: the definitions of the four variables make it true. The equation is useful because it shows that if one of the variables changes, one or more of the others must also change to maintain the equality. For example, if the quantity of money increases and the velocity of money stays unchanged, then either the price or the number of transactions must rise.

From Transactions to Income

When studying the role of money in the economy, economists usually use a slightly different version of the quantity equation than the one just introduced. The problem with the first equation is that the number of transactions is difficult to measure. To solve this problem, the number of transactions T is replaced by the total output of the economy Y .

Transactions and output are related, because the more the economy produces, the more goods are bought and sold. They are not the same, however. When one person sells a used car to another person, for example, they make a transaction using money, even though the used car is not part of current output. Nonetheless, the dollar value of transactions is roughly proportional to the dollar value of output.

If Y denotes the amount of output and P denotes the price of one unit of output, then the dollar value of output is PY . We encountered measures for these variables when we discussed the national income accounts in Chapter 2: Y is real GDP, P is the GDP deflator, and PY is nominal GDP. The quantity equation becomes

$$\begin{array}{ccccccc} \text{Money} & \times & \text{Velocity} & = & \text{Price} & \times & \text{Output} \\ M & \times & V & = & P & \times & Y. \end{array}$$

Because Y is also total income, V in this version of the quantity equation is called the **income velocity of money**. The income velocity of money tells us the number of times a dollar bill enters someone's income in a given period of time. This version of the quantity equation is the most common, and it is the one we use from now on.

The Money Demand Function and the Quantity Equation

When we analyze how money affects the economy, it is often useful to express the quantity of money in terms of the quantity of goods and services it can buy. This amount, M/P , is called **real money balances**.

Real money balances measure the purchasing power of the stock of money. For example, consider an economy that produces only bread. If the quantity of money is \$10, and the price of a loaf is \$0.50, then real money balances are 20 loaves of bread. That is, at current prices, the stock of money in the economy is able to buy 20 loaves.

A **money demand function** is an equation that shows what determines the quantity of real money balances people wish to hold. A simple money demand function is

$$(M/P)^d = kY,$$

where k is a constant that tells us how much money people want to hold for every dollar of income. This equation states that the quantity of real money balances demanded is proportional to real income.

The money demand function is like the demand function for a particular good. Here the “good” is the convenience of holding real money balances. Just as owning an automobile makes it easier for a person to travel, holding money makes it easier to make transactions. Therefore, just as higher income leads to a greater demand for automobiles, higher income also leads to a greater demand for real money balances.

This money demand function offers another way to view the quantity equation. To see this, add to the money demand function the condition that the demand for real money balances $(M/P)^d$ must equal the supply M/P . Therefore,

$$M/P = kY.$$

A simple rearrangement of terms changes this equation into

$$M(1/k) = PY,$$

which can be written as

$$MV = PY,$$

where $V = 1/k$. This simple mathematics shows the link between the demand for money and the velocity of money. When people want to hold a lot of money for each dollar of income (k is large), money changes hands infrequently (V is small). Conversely, when people want to hold only a little money (k is small), money changes hands frequently (V is large). In other words, the money demand parameter k and the velocity of money V are opposite sides of the same coin.

The Assumption of Constant Velocity

The quantity equation can be viewed as a definition: it defines velocity V as the ratio of nominal GDP, PY , to the quantity of money M . Yet if we make the additional assumption that the velocity of money is constant, then the quantity equation becomes a useful theory of the effects of money, called the **quantity theory of money**.

As with many of the assumptions in economics, the assumption of constant velocity is only an approximation to reality. Velocity does change if the money demand function changes. For example, when automatic teller machines were introduced, people could reduce their average money holdings, which meant a fall in the money demand parameter k and an increase in velocity V . Nonetheless, experience shows that the assumption of constant velocity provides a good approximation in many situations. Let's therefore assume that velocity is constant and see what this assumption implies about the effects of the money supply on the economy.

Once we assume that velocity is constant, the quantity equation can be seen as a theory of what determines nominal GDP. The quantity equation says

$$M\bar{V} = PY,$$

where the bar over V means that velocity is fixed. Therefore, a change in the quantity of money (M) must cause a proportionate change in nominal GDP (PY). That is, if velocity is fixed, the quantity of money determines the dollar value of the economy's output.

Money, Prices, and Inflation

We now have a theory to explain what determines the economy's overall level of prices. The theory has three building blocks:

1. The factors of production and the production function determine the level of output Y . We borrow this conclusion from Chapter 3.
2. The money supply determines the nominal value of output, PY . This conclusion follows from the quantity equation and the assumption that the velocity of money is fixed.
3. The price level P is then the ratio of the nominal value of output, PY , to the level of output Y .

In other words, the productive capability of the economy determines real GDP, the quantity of money determines nominal GDP, and the GDP deflator is the ratio of nominal GDP to real GDP.

This theory explains what happens when the Fed changes the supply of money. Because velocity is fixed, any change in the supply of money leads to a proportionate change in nominal GDP. Because the factors of production and the production function have already determined real GDP, the change in nominal GDP must represent a change in the price level. Hence, the quantity theory implies that the price level is proportional to the money supply.

Because the inflation rate is the percentage change in the price level, this theory of the price level is also a theory of the inflation rate. The quantity equation, written in percentage-change form, is

$$\% \text{ Change in } M + \% \text{ Change in } V = \% \text{ Change in } P + \% \text{ Change in } Y.$$

Consider each of these four terms. First, the percentage change in the quantity of money M is under the control of the central bank. Second, the percentage change in velocity V reflects shifts in money demand; we have assumed that velocity is constant, so the percentage change in velocity is zero. Third, the percentage change in the price level P is the rate of inflation; this is the variable in the equation that we would like to explain. Fourth, the percentage change in output Y depends on growth in the factors of production and on technological progress, which for our present purposes we can take as given. This analysis tells us that (except for a constant that depends on exogenous growth in output) the growth in the money supply determines the rate of inflation.

Thus, the quantity theory of money states that the central bank, which controls the money supply, has ultimate control over the rate of inflation. If the central bank keeps the money supply stable, the price level will be stable. If the central bank increases the money supply rapidly, the price level will rise rapidly.


CASE STUDY

Inflation and Money Growth

“Inflation is always and everywhere a monetary phenomenon.” So wrote Milton Friedman, the great economist who won the Nobel Prize in economics in 1976. The quantity theory of money leads us to agree that the growth in the quantity of money is the primary determinant of the inflation rate. Yet Friedman’s claim is empirical, not theoretical. To evaluate his claim, and to judge the usefulness of our theory, we need to look at data on money and prices.

Friedman, together with fellow economist Anna Schwartz, wrote two treatises on monetary history that documented the sources and effects of changes in the quantity of money over the past century.³ Figure 4-1 uses some of their data and

figure 4-1




Historical Data on U.S. Inflation and Money Growth In this scatterplot of money growth and inflation, each point represents a decade. The horizontal axis shows the average growth in the money supply (as measured by M2) over the decade, and the vertical axis shows the average rate of inflation (as measured by the GDP deflator). The positive correlation between money growth and inflation is evidence for the quantity theory’s prediction that high money growth leads to high inflation.

Source: For the data through the 1960s: Milton Friedman and Anna J. Schwartz, *Monetary Trends in the United States and the United Kingdom: Their Relation to Income, Prices, and Interest Rates, 1867–1975* (Chicago: University of Chicago Press, 1982). For recent data: U.S. Department of Commerce, Federal Reserve Board.

³ Milton Friedman and Anna J. Schwartz, *A Monetary History of the United States, 1867–1960* (Princeton, NJ: Princeton University Press, 1963); Milton Friedman and Anna J. Schwartz, *Monetary Trends in the United States and the United Kingdom: Their Relation to Income, Prices, and Interest Rates, 1867–1975* (Chicago: University of Chicago Press, 1982).

plots the average rate of money growth and the average rate of inflation in the United States over each decade since the 1870s. The data verify the link between inflation and growth in the quantity of money. Decades with high money growth tend to have high inflation, and decades with low money growth tend to have low inflation.

figure 4-2



International Data on Inflation and Money Growth In this scatterplot, each point represents a country. The horizontal axis shows the average growth in the money supply (as measured by currency plus demand deposits) during the 1990s, and the vertical axis shows the average rate of inflation (as measured by the GDP deflator). Once again, the positive correlation is evidence for the quantity theory's prediction that high money growth leads to high inflation.

Source: International Financial Statistics.

Figure 4-2 examines the same question with international data. It shows the average rate of inflation and the average rate of money growth in more than 100 countries during the 1990s. Again, the link between money growth and inflation is clear. Countries with high money growth tend to have high inflation, and countries with low money growth tend to have low inflation.

If we looked at monthly data on money growth and inflation, rather than data for 10-year periods, we would not see as close a connection between these two variables. This theory of inflation works best in the long run, not in the short run. We examine the short-run impact of changes in the quantity of money when we turn to economic fluctuations in Part IV of this book.

4-3 Seigniorage: The Revenue From Printing Money

So far, we have seen how growth in the money supply causes inflation. But what might ever induce the government to increase the money supply? Here we examine one answer to this question.

Let's start with an indisputable fact: all governments spend money. Some of this spending is to buy goods and services (such as roads and police), and some is to provide transfer payments (for the poor and elderly, for example). A government can finance its spending in three ways. First, it can raise revenue through taxes, such as personal and corporate income taxes. Second, it can borrow from the public by selling government bonds. Third, it can print money.

The revenue raised through the printing of money is called **seigniorage**. The term comes from *seigneur*, the French word for “feudal lord.” In the Middle Ages, the lord had the exclusive right on his manor to coin money. Today this right belongs to the central government, and it is one source of revenue.

When the government prints money to finance expenditure, it increases the money supply. The increase in the money supply, in turn, causes inflation. Printing money to raise revenue is like imposing an *inflation tax*.

At first it may not be obvious that inflation can be viewed as a tax. After all, no one receives a bill for this tax—the government merely prints the money it needs. Who then pays the inflation tax? The answer is the holders of money. As prices rise, the real value of the money in your wallet falls. When the government prints new money for its use, it makes the old money in the hands of the public less valuable. Thus, inflation is like a tax on holding money.

The amount raised by printing money varies from country to country. In the United States, the amount has been small: seigniorage has usually accounted for less than 3 percent of government revenue. In Italy and Greece, seigniorage has often been more than 10 percent of government revenue.⁴ In countries experiencing hyperinflation, seigniorage is often the government's chief source of revenue—indeed, the need to print money to finance expenditure is a primary cause of hyperinflation.

CASE STUDY

Paying for the American Revolution

Although seigniorage has not been a major source of revenue for the U.S. government in recent history, the situation was very different two centuries ago. Beginning in 1775 the Continental Congress needed to find a way to finance the Revolution, but it had limited ability to raise revenue through taxation. It, therefore, relied on the printing of fiat money to help pay for the war.

The Continental Congress's reliance on seigniorage increased over time. In 1775 new issues of continental currency were about \$6 million. This amount

⁴ Stanley Fischer, “Seigniorage and the Case for a National Money,” *Journal of Political Economy* 90 (April 1982): 295–313.

increased to \$19 million in 1776, \$13 million in 1777, \$63 million in 1778, and \$125 million in 1779.

Not surprisingly, this rapid growth in the money supply led to massive inflation. At the end of the war, the price of gold measured in continental dollars was more than 100 times its level of only a few years earlier. The large quantity of the continental currency made the continental dollar nearly worthless. This experience also gave birth to a once-popular expression: people used to say something was “not worth a continental,” meaning that the item had little real value.

4-4 Inflation and Interest Rates

As we first discussed in Chapter 3, interest rates are among the most important macroeconomic variables. In essence, they are the prices that link the present and the future. Here we discuss the relationship between inflation and interest rates.

Two Interest Rates: Real and Nominal

Suppose you deposit your savings in a bank account that pays 8 percent interest annually. Next year, you withdraw your savings and the accumulated interest. Are you 8 percent richer than you were when you made the deposit a year earlier?

The answer depends on what “richer” means. Certainly, you have 8 percent more dollars than you had before. But if prices have risen, so that each dollar buys less, then your purchasing power has not risen by 8 percent. If the inflation rate was 5 percent, then the amount of goods you can buy has increased by only 3 percent. And if the inflation rate was 10 percent, then your purchasing power has fallen by 2 percent.

Economists call the interest rate that the bank pays the **nominal interest rate** and the increase in your purchasing power the **real interest rate**. If i denotes the nominal interest rate, r the real interest rate, and π the rate of inflation, then the relationship among these three variables can be written as

$$r = i - \pi.$$

The real interest rate is the difference between the nominal interest rate and the rate of inflation.⁵

The Fisher Effect

Rearranging the terms in our equation for the real interest rate, we can show that the nominal interest rate is the sum of the real interest rate and the inflation rate:

$$i = r + \pi.$$

⁵ *Mathematical note:* This equation relating the real interest rate, nominal interest rate, and inflation rate is only an approximation. The exact formula is $(1 + r) = (1 + i)/(1 + \pi)$. The approximation in the text is reasonably accurate as long as r , i , and π are relatively small (say, less than 20 percent per year).

The equation written in this way is called the **Fisher equation**, after economist Irving Fisher (1867–1947). It shows that the nominal interest rate can change for two reasons: because the real interest rate changes or because the inflation rate changes.

Once we separate the nominal interest rate into these two parts, we can use this equation to develop a theory that explains the nominal interest rate. Chapter 3 showed that the real interest rate adjusts to equilibrate saving and investment. The quantity theory of money shows that the rate of money growth determines the rate of inflation. The Fisher equation then tells us to add the real interest rate and the inflation rate together to determine the nominal interest rate.


The quantity theory and the Fisher equation together tell us how money growth affects the nominal interest rate. *According to the quantity theory, an increase in the rate of money growth of 1 percent causes a 1-percent increase in the rate of inflation. According to the Fisher equation, a 1-percent increase in the rate of inflation in turn causes a 1-percent increase in the nominal interest rate.* The one-for-one relation between the inflation rate and the nominal interest rate is called the **Fisher effect**.

CASE STUDY

Inflation and Nominal Interest Rates

How useful is the Fisher effect in explaining interest rates? To answer this question we look at two types of data on inflation and nominal interest rates.

figure 4-3




Inflation and Nominal Interest Rates Over Time This figure plots the nominal interest rate (on three-month Treasury bills) and the inflation rate (as measured by the CPI) in the United States since 1954. It shows the Fisher effect: higher inflation leads to a higher nominal interest rate.

Source: Federal Reserve and U.S. Department of Labor.

Figure 4-3 shows the variation over time in the nominal interest rate and the inflation rate in the United States. You can see that the Fisher effect has done a good job explaining fluctuations in the nominal interest rate over the past half century. When inflation is high, nominal interest rates are typically high, and when inflation is low, nominal interest rates are typically low as well.

figure 4-4



Inflation and Nominal Interest Rates Across Countries This scatterplot exhibits the average nominal interest rate on short-term Treasury bills and the average inflation rate in 77 countries during the 1990s. The positive correlation between the inflation rate and the nominal interest rate is evidence for the Fisher effect.

Source: International Financial Statistics.

Similar support for the Fisher effect comes from examining the variation across countries. As Figure 4-4 shows, a nation's inflation rate and its nominal interest rate are related. Countries with high inflation tend to have high nominal interest rates as well, and countries with low inflation tend to have low nominal interest rates.

The link between inflation and interest rates is well known to Wall Street investment firms. Because bond prices move inversely with interest rates, one can get rich by predicting correctly the direction in which interest rates will move. Many Wall Street firms hire *Fed watchers* to monitor monetary policy and news about inflation in order to anticipate changes in interest rates.

Two Real Interest Rates: *Ex Ante* and *Ex Post*

When a borrower and lender agree on a nominal interest rate, they do not know what the inflation rate over the term of the loan will be. Therefore, we must distinguish between two concepts of the real interest rate: the real interest rate the borrower and lender expect when the loan is made, called the ***ex ante real interest rate***, and the real interest rate actually realized, called the ***ex post real interest rate***.

Although borrowers and lenders cannot predict future inflation with certainty, they do have some expectation of the inflation rate. Let π denote actual future inflation and π^e the expectation of future inflation. The *ex ante* real interest rate is $i - \pi^e$, and the *ex post* real interest rate is $i - \pi$. The two real interest rates differ when actual inflation π differs from expected inflation π^e .

How does this distinction between actual and expected inflation modify the Fisher effect? Clearly, the nominal interest rate cannot adjust to actual inflation, because actual inflation is not known when the nominal interest rate is set. The nominal interest rate can adjust only to expected inflation. The Fisher effect is more precisely written as

$$i = r + \pi^e.$$

The *ex ante* real interest rate r is determined by equilibrium in the market for goods and services, as described by the model in Chapter 3. The nominal interest rate i moves one-for-one with changes in expected inflation π^e .

CASE STUDY

Nominal Interest Rates in the Nineteenth Century

Although recent data show a positive relationship between nominal interest rates and inflation rates, this finding is not universal. In data from the late nineteenth and early twentieth centuries, high nominal interest rates did not accompany high inflation. The apparent absence of any Fisher effect during this time puzzled Irving Fisher. He suggested that inflation “caught merchants napping.”

How should we interpret the absence of an apparent Fisher effect in nineteenth-century data? Does this period of history provide evidence against the adjustment of nominal interest rates to inflation? Recent research suggests that this period has little to tell us about the validity of the Fisher effect. The reason is that the Fisher effect relates the nominal interest rate to expected inflation and, according to this research, inflation at this time was largely unexpected.

Although expectations are not observable, we can draw inferences about them by examining the persistence of inflation. In recent experience, inflation has been very persistent: when it is high one year, it tends to be high the next year as well. Therefore, when people have observed high inflation, it has been rational for them to expect high inflation in the future. By contrast, during the nineteenth century, when the gold standard was in effect, inflation had

little persistence. High inflation in one year was just as likely to be followed the next year by low inflation as by high inflation. Therefore, high inflation did not imply high expected inflation and did not lead to high nominal interest rates. So, in a sense, Fisher was right to say that inflation “caught merchants napping.”⁶

4-5 The Nominal Interest Rate and the Demand for Money

The quantity theory is based on a simple money demand function: it assumes that the demand for real money balances is proportional to income. Although the quantity theory is a good place to start when analyzing the effects of money on the economy, it is not the whole story. Here we add another determinant of the quantity of money demanded—the nominal interest rate.

The Cost of Holding Money

The money you hold in your wallet does not earn interest. If, instead of holding that money, you used it to buy government bonds or deposited it in a savings account, you would earn the nominal interest rate. The nominal interest rate is the opportunity cost of holding money: it is what you give up by holding money rather than bonds.

Another way to see that the cost of holding money equals the nominal interest rate is by comparing the real returns on alternative assets. Assets other than money, such as government bonds, earn the real return r . Money earns an expected real return of $-\pi^e$, because its real value declines at the rate of inflation. When you hold money, you give up the difference between these two returns. Thus, the cost of holding money is $r - (-\pi^e)$, which the Fisher equation tells us is the nominal interest rate i .

Just as the quantity of bread demanded depends on the price of bread, the quantity of money demanded depends on the price of holding money. Hence, the demand for real money balances depends both on the level of income and on the nominal interest rate. We write the general money demand function as

$$(M/P)^d = L(i, Y).$$

The letter L is used to denote money demand because money is the economy’s most liquid asset (the asset most easily used to make transactions). This equation states that the demand for the liquidity of real money balances is a function of

⁶ Robert B. Barsky, “The Fisher Effect and the Forecastability and Persistence of Inflation,” *Journal of Monetary Economics* 19 (January 1987): 3–24.

income and the nominal interest rate. The higher the level of income Y , the greater the demand for real money balances. The higher the nominal interest rate i , the lower the demand for real money balances.

Future Money and Current Prices

Money, prices, and interest rates are now related in several ways. Figure 4-5 illustrates the linkages we have discussed. As the quantity theory of money explains, money supply and money demand together determine the equilibrium price level. Changes in the price level are, by definition, the rate of inflation. Inflation, in turn, affects the nominal interest rate through the Fisher effect. But now, because the nominal interest rate is the cost of holding money, the nominal interest rate feeds back to affect the demand for money.

Consider how the introduction of this last link affects our theory of the price level. First, equate the supply of real money balances M/P to the demand $L(i, Y)$:


$$M/P = L(i, Y).$$

Next, use the Fisher equation to write the nominal interest rate as the sum of the real interest rate and expected inflation:

$$M/P = L(r + \pi^e, Y).$$

This equation states that the level of real money balances depends on the expected rate of inflation.

figure 4-5



The Linkages Among Money, Prices, and Interest Rates This figure illustrates the relationships among money, prices, and interest rates. Money supply and money demand determine the price level. Changes in the price level determine the inflation rate. The inflation rate influences the nominal interest rate. Because the nominal interest rate is the cost of holding money, it may affect money demand. This last link (shown as a blue line) is omitted from the basic quantity theory of money.

The last equation tells a more sophisticated story than the quantity theory about the determination of the price level. The quantity theory of money says that today's money supply determines today's price level. This conclusion remains partly true: if the nominal interest rate and the level of output are held constant, the price level moves proportionately with the money supply. Yet the nominal interest rate is not constant; it depends on expected inflation, which in turn depends on growth in the money supply. The presence of the nominal interest rate in the money demand function yields an additional channel through which money supply affects the price level.

This general money demand equation implies that the price level depends not only on today's money supply but also on the money supply expected in the future. To see why, suppose the Fed announces that it will raise the money supply in the future, but it does not change the money supply today. This announcement causes people to expect higher money growth and higher inflation. Through the Fisher effect, this increase in expected inflation raises the nominal interest rate. The higher nominal interest rate reduces the demand for real money balances. Because the quantity of money has not changed, the reduced demand for real money balances leads to a higher price level. Hence, higher expected money growth in the future leads to a higher price level today.

The effect of money on prices is complex. The appendix to this chapter works out the mathematics relating the price level to current and future money. The conclusion of the analysis is that the price level depends on a weighted average of the current money supply and the money supply expected to prevail in the future.

4-6 The Social Costs of Inflation

Our discussion of the causes and effects of inflation does not tell us much about the social problems that result from inflation. We turn to those problems now.

The Layman's View and the Classical Response

If you ask the average person why inflation is a social problem, he will probably answer that inflation makes him poorer. "Each year my boss gives me a raise, but prices go up and that takes some of my raise away from me." The implicit assumption in this statement is that if there were no inflation, he would get the same raise and be able to buy more goods.

This complaint about inflation is a common fallacy. As we know from Chapter 3, the purchasing power of labor—the real wage—depends on the marginal productivity of labor, not on how much money the government chooses to print. If the government reduced inflation by slowing the rate of money growth, workers would not see their real wage increasing more rapidly. Instead, when

inflation slowed, firms would increase the prices of their products less each year and, as a result, would give their workers smaller raises.

According to the classical theory of money, a change in the overall price level is like a change in the units of measurement. It is as if we switched from measuring distances in feet to measuring them in inches: numbers get larger, but nothing really changes. Imagine that tomorrow morning you wake up and find that, for some reason, all dollar figures in the economy have been multiplied by ten. The price of everything you buy has increased tenfold, but so has your wage and the value of your savings. What difference would this make? All numbers would have an extra zero at the end, but nothing else would change. Your economic well-being depends on relative prices, not the overall price level.

Why, then, is a persistent increase in the price level a social problem? It turns out that the costs of inflation are subtle. Indeed, economists disagree about the size of the social costs. To the surprise of many laymen, some economists argue that the costs of inflation are small—at least for the moderate rates of inflation that most countries have experienced in recent years.⁷

CASE STUDY

What Economists and the Public Say About Inflation

As we have been discussing, laymen and economists hold very different views about the costs of inflation. Economist Robert Shiller has documented this difference of opinion in a survey of the two groups. The survey results are striking, for they show how the study of economics changes a person's attitudes.

In one question, Shiller asked people whether their “biggest gripe about inflation” was that “inflation hurts my real buying power, it makes me poorer.” Of the general public, 77 percent agreed with this statement, compared to only 12 percent of economists. Shiller also asked people whether they agreed with the following statement: “When I see projections about how many times more a college education will cost, or how many times more the cost of living will be in coming decades, I feel a sense of uneasiness; these inflation projections really make me worry that my own income will not rise as much as such costs will.” Among the general public, 66 percent said they fully agreed with this statement, while only 5 percent of economists agreed with it.

Survey respondents were asked to judge the seriousness of inflation as a policy problem: “Do you agree that preventing high inflation is an important national priority, as important as preventing drug abuse or preventing deterioration in the quality of our schools?” Fifty-two percent of laymen, but only 18 percent of economists, fully agreed with this view. Apparently, inflation worries the public much more than it does the economics profession.

⁷ See, for example, Chapter 2 of Alan Blinder, *Hard Heads, Soft Hearts: Tough-Minded Economics for a Just Society* (Reading, MA: Addison Wesley, 1987).

The public's distaste for inflation may be psychological. Shiller asked those surveyed if they agreed with the following statement: "I think that if my pay went up I would feel more satisfaction in my job, more sense of fulfillment, even if prices went up just as much." Of the public, 49 percent fully or partly agreed with this statement, compared to 8 percent of economists.

Do these survey results mean that laymen are wrong and economists are right about the costs of inflation? Not necessarily. But economists do have the advantage of having given the issue more thought. So let's now consider what some of the costs of inflation might be.⁸

The Costs of Expected Inflation

Consider first the case of expected inflation. Suppose that every month the price level rose by 1 percent. What would be the social costs of such a steady and predictable 12-percent annual inflation?

One cost is the distortion of the inflation tax on the amount of money people hold. As we have already discussed, a higher inflation rate leads to a higher nominal interest rate, which in turn leads to lower real money balances. If people are to hold lower money balances on average, they must make more frequent trips to the bank to withdraw money—for example, they might withdraw \$50 twice a week rather than \$100 once a week. The inconvenience of reducing money holding is metaphorically called the **shoe-leather cost** of inflation, because walking to the bank more often causes one's shoes to wear out more quickly.

A second cost of inflation arises because high inflation induces firms to change their posted prices more often. Changing prices is sometimes costly: for example, it may require printing and distributing a new catalog. These costs are called **menu costs**, because the higher the rate of inflation, the more often restaurants have to print new menus.

A third cost of inflation arises because firms facing menu costs change prices infrequently; therefore, the higher the rate of inflation, the greater the variability in relative prices. For example, suppose a firm issues a new catalog every January. If there is no inflation, then the firm's prices relative to the overall price level are constant over the year. Yet if inflation is 1 percent per month, then from the beginning to the end of the year the firm's relative prices fall by 12 percent. Sales from this catalog will tend to be low early in the year (when its prices are relatively high) and high later in the year (when its prices are relatively low). Hence, when inflation induces variability in relative prices, it leads to microeconomic inefficiencies in the allocation of resources.

⁸ Robert J. Shiller, "Why Do People Dislike Inflation?" in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997).

A fourth cost of inflation results from the tax laws. Many provisions of the tax code do not take into account the effects of inflation. Inflation can alter individuals' tax liability, often in ways that lawmakers did not intend.

One example of the failure of the tax code to deal with inflation is the tax treatment of capital gains. Suppose you buy some stock today and sell it a year from now at the same real price. It would seem reasonable for the government not to levy a tax, because you have earned no real income from this investment. Indeed, if there is no inflation, a zero tax liability would be the outcome. But suppose the rate of inflation is 12 percent and you initially paid \$100 per share for the stock; for the real price to be the same a year later, you must sell the stock for \$112 per share. In this case the tax code, which ignores the effects of inflation, says that you have earned \$12 per share in income, and the government taxes you on this capital gain. The problem, of course, is that the tax code measures income as the nominal rather than the real capital gain. In this example, and in many others, inflation distorts how taxes are levied.

A fifth cost of inflation is the inconvenience of living in a world with a changing price level. Money is the yardstick with which we measure economic transactions. When there is inflation, that yardstick is changing in length. To continue the analogy, suppose that Congress passed a law specifying that a yard would equal 36 inches in 2002, 35 inches in 2003, 34 inches in 2004, and so on. Although the law would result in no ambiguity, it would be highly inconvenient. When someone measured a distance in yards, it would be necessary to specify whether the measurement was in 2002 yards or 2003 yards; to compare distances measured in different years, one would need to make an "inflation" correction. Similarly, the dollar is a less useful measure when its value is always changing.

For example, a changing price level complicates personal financial planning. One important decision that all households face is how much of their income to consume today and how much to save for retirement. A dollar saved today and invested at a fixed nominal interest rate will yield a fixed dollar amount in the future. Yet the real value of that dollar amount—which will determine the retiree's living standard—depends on the future price level. Deciding how much to save would be much simpler if people could count on the price level in 30 years being similar to its level today.

The Costs of Unexpected Inflation

Unexpected inflation has an effect that is more pernicious than any of the costs of steady, anticipated inflation: it arbitrarily redistributes wealth among individuals. You can see how this works by examining long-term loans. Most loan agreements specify a nominal interest rate, which is based on the rate of inflation expected at the time of the agreement. If inflation turns out differently from what was expected, the *ex post* real return that the debtor pays to the creditor differs from what both parties anticipated. On the one hand, if inflation turns out to be higher than expected, the debtor wins and the creditor loses because the debtor repays the loan with less valuable dollars. On the other hand, if inflation turns out to be lower than expected, the creditor wins and the debtor loses because the repayment is worth more than the two parties anticipated.

Consider, for example, a person taking out a mortgage in 1960. At the time, a 30-year mortgage had an interest rate of about 6 percent per year. This rate was based on a low rate of expected inflation—inflation over the previous decade had averaged only 2.5 percent. The creditor probably expected to receive a real return of about 3.5 percent, and the debtor expected to pay this real return. In fact, over the life of the mortgage, the inflation rate averaged 5 percent, so the *ex post* real return was only 1 percent. This unanticipated inflation benefited the debtor at the expense of the creditor.

Unanticipated inflation also hurts individuals on fixed pensions. Workers and firms often agree on a fixed nominal pension when the worker retires (or even earlier). Because the pension is deferred earnings, the worker is essentially providing the firm a loan: the worker provides labor services to the firm while young but does not get fully paid until old age. Like any creditor, the worker is hurt when inflation is higher than anticipated. Like any debtor, the firm is hurt when inflation is lower than anticipated.

These situations provide a clear argument against variable inflation. The more variable the rate of inflation, the greater the uncertainty that both debtors and creditors face. Because most people are *risk averse*—they dislike uncertainty—the unpredictability caused by highly variable inflation hurts almost everyone.

Given these effects of uncertain inflation, it is puzzling that nominal contracts are so prevalent. One might expect debtors and creditors to protect themselves from this uncertainty by writing contracts in real terms—that is, by indexing to some measure of the price level. In economies with high and variable inflation, indexation is often widespread; sometimes this indexation takes the form of writing contracts using a more stable foreign currency. In economies with moderate inflation, such as the United States, indexation is less common. Yet even in the United States, some long-term obligations are indexed. For example, Social Security benefits for the elderly are adjusted annually in response to changes in the consumer price index. And in 1997, the U.S. federal government issued inflation-indexed bonds for the first time.

Finally, in thinking about the costs of inflation, it is important to note a widely documented but little understood fact: high inflation is variable inflation. That is, countries with high average inflation also tend to have inflation rates that change greatly from year to year. The implication is that if a country decides to pursue a high-inflation monetary policy, it will likely have to accept highly variable inflation as well. As we have just discussed, highly variable inflation increases uncertainty for both creditors and debtors by subjecting them to arbitrary and potentially large redistributions of wealth.

CASE STUDY

The Free Silver Movement, the Election of 1896, and the Wizard of Oz

The redistributions of wealth caused by unexpected changes in the price level are often a source of political turmoil, as evidenced by the Free Silver movement in the late nineteenth century. From 1880 to 1896 the price level in the

United States fell 23 percent. This deflation was good for creditors, primarily the bankers of the Northeast, but it was bad for debtors, primarily the farmers of the South and West. One proposed solution to this problem was to replace the gold standard with a bimetallic standard, under which both gold and silver could be minted into coin. The move to a bimetallic standard would increase the money supply and stop the deflation.

The silver issue dominated the presidential election of 1896. William McKinley, the Republican nominee, campaigned on a platform of preserving the gold standard. William Jennings Bryan, the Democratic nominee, supported the bimetallic standard. In a famous speech, Bryan proclaimed, “You shall not press down upon the brow of labor this crown of thorns, you shall not crucify mankind upon a cross of gold.” Not surprisingly, McKinley was the candidate of the conservative eastern establishment, while Bryan was the candidate of the southern and western populists.

This debate over silver found its most memorable expression in a children’s book, *The Wizard of Oz*. Written by a midwestern journalist, L. Frank Baum, just after the 1896 election, it tells the story of Dorothy, a girl lost in a strange land far from her home in Kansas. Dorothy (representing traditional American values) makes three friends: a scarecrow (the farmer), a tin woodman (the industrial worker), and a lion whose roar exceeds his might (William Jennings Bryan). Together, the four of them make their way along a perilous yellow brick road (the gold standard), hoping to find the Wizard who will help Dorothy return home. Eventually they arrive in Oz (Washington), where everyone sees the world through green glasses (money). The Wizard (William McKinley) tries to be all things to all people but turns out to be a fraud. Dorothy’s problem is solved only when she learns about the magical power of her silver slippers.⁹

Although the Republicans won the election of 1896 and the United States stayed on a gold standard, the Free Silver advocates got the inflation that they wanted. Around the time of the election, gold was discovered in Alaska, Australia, and South Africa. In addition, gold refiners devised the cyanide process, which facilitated the extraction of gold from ore. These developments led to increases in the money supply and in prices. From 1896 to 1910 the price level rose 35 percent.

One Benefit of Inflation

So far, we have discussed the many costs of inflation. These costs lead many economists to conclude that monetary policymakers should aim for zero inflation. Yet there is another side to the story. Some economists believe that a little bit of inflation—say, 2 or 3 percent per year—can be a good thing.

⁹ The movie made 40 years later hid much of the allegory by changing Dorothy’s slippers from silver to ruby. For more on this topic, see Henry M. Littlefield, “The Wizard of Oz: Parable on Populism,” *American Quarterly* 16 (Spring 1964): 47–58; and Hugh Rockoff, “The Wizard of Oz as a Monetary Allegory,” *Journal of Political Economy* 98 (August 1990): 739–760.



Keynes (and Lenin) on the Cost of Inflation

The great economist John Maynard Keynes was no friend of inflation, as this chapter's opening quotation indicates. Here is the more complete passage from his famous book, *The Economic Consequences of the Peace*, in which Keynes predicted (correctly) that the treaty imposed on Germany after World War I would lead to economic hardship and renewed international tensions:

Lenin is said to have declared that the best way to destroy the Capitalist System was to debauch the currency. By a continuing process of inflation, governments can confiscate, secretly and unobserved, an important part of the wealth of their citizens. By this method they not only confiscate, but they confiscate *arbitrarily*; and, while the process impoverishes many, it actually enriches some. The sight of this arbitrary rearrangement of riches strikes not only at security, but at confidence in the equity of the existing distribution of wealth. Those to whom the system brings windfalls, beyond their deserts and even beyond their expectations or desires, become "profiteers," who are the object of the hatred of the bourgeoisie,

whom the inflationism has impoverished, not less than of the proletariat. As the inflation proceeds and the real value of the currency fluctuates wildly from month to month, all permanent relations between debtors and creditors, which form the ultimate foundation of capitalism, become so utterly disordered as to be almost meaningless; and the process of wealth-getting degenerates into a gamble and a lottery.

Lenin was certainly right. There is no subtler, no surer means of overturning the existing basis of society than to debauch the currency. The process engages all the hidden forces of economic law on the side of destruction, and does it in a manner which not one man in a million is able to diagnose.¹⁰

History has given ample support to this assessment. A recent example occurred in Russia in 1998, where many citizens saw high rates of inflation wipe out their ruble-denominated savings. And, as Lenin would have predicted, this inflation put the country's burgeoning capitalist system in serious jeopardy.

The argument for moderate inflation starts with the observation that cuts in nominal wages are rare: firms are reluctant to cut their workers' nominal wages, and workers are reluctant to accept such cuts. A 2-percent wage cut in a zero-inflation world is, in real terms, the same as a 3-percent raise with 5-percent inflation, but workers do not always see it that way. The 2-percent wage cut may seem like an insult, whereas the 3-percent raise is, after all, still a raise. Empirical studies confirm that nominal wages rarely fall.

This fact suggests that some inflation may make labor markets work better. The supply and demand for different kinds of labor are always changing. Sometimes an increase in supply or decrease in demand leads to a fall in the equilibrium real wage for a group of workers. If nominal wages can't be cut, then the only way to cut real wages is to allow inflation to do the job. Without inflation, the real wage will be stuck above the equilibrium level, resulting in higher unemployment.

For this reason, some economists argue that inflation "greases the wheels" of labor markets. Only a little inflation is needed: an inflation rate of 2 percent

¹⁰ John Maynard Keynes, *The Economic Consequences of the Peace* (London: Macmillan, 1920): 219–220.

lets real wages fall by 2 percent per year, or 20 percent per decade, without cuts in nominal wages. Such automatic reductions in real wages are impossible with zero inflation.¹¹

4-7 Hyperinflation

Hyperinflation is often defined as inflation that exceeds 50 percent per month, which is just over 1 percent per day. Compounded over many months, this rate of inflation leads to very large increases in the price level. An inflation rate of 50 percent per month implies a more than 100-fold increase in the price level over a year, and a more than 2-million-fold increase over three years. Here we consider the costs and causes of such extreme inflation.

The Costs of Hyperinflation

Although economists debate whether the costs of moderate inflation are large or small, no one doubts that hyperinflation extracts a high toll on society. The costs are qualitatively the same as those we discussed earlier. When inflation reaches extreme levels, however, these costs are more apparent because they are so severe.

The shoeleather costs associated with reduced money holding, for instance, are serious under hyperinflation. Business executives devote much time and energy to cash management when cash loses its value quickly. By diverting this time and energy from more socially valuable activities, such as production and investment decisions, hyperinflation makes the economy run less efficiently.

Menu costs also become larger under hyperinflation. Firms have to change prices so often that normal business practices, such as printing and distributing catalogs with fixed prices, become impossible. In one restaurant during the German hyperinflation of the 1920s, a waiter would stand up on a table every 30 minutes to call out the new prices.

Similarly, relative prices do not do a good job of reflecting true scarcity during hyperinflations. When prices change frequently by large amounts, it is hard for customers to shop around for the best price. Highly volatile and rapidly rising prices can alter behavior in many ways. According to one report, when patrons entered a pub during the German hyperinflation, they would often buy two pitchers of beer. Although the second pitcher would lose value by getting warm over time, it would lose value less rapidly than the money left sitting in the patron's wallet.

¹¹ For a recent paper examining this benefit of inflation, see George A. Akerlof, William T. Dickens, and George L. Perry, "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity*, 1996:1, pp. 1–76.

Tax systems are also distorted by hyperinflation—but in ways that are quite different than under moderate inflation. In most tax systems there is a delay between the time a tax is levied and the time the tax is paid to the government. In the United States, for example, taxpayers are required to make estimated income tax payments every three months. This short delay does not matter much under low inflation. By contrast, during hyperinflation, even a short delay greatly reduces real tax revenue. By the time the government gets the money it is due, the money has fallen in value. As a result, once hyperinflations start, the real tax revenue of the government often falls substantially.

Finally, no one should underestimate the sheer inconvenience of living with hyperinflation. When carrying money to the grocery store is as burdensome as carrying the groceries back home, the monetary system is not doing its best to facilitate exchange. The government tries to overcome this problem by adding more and more zeros to the paper currency, but often it cannot keep up with the exploding price level.

Eventually, these costs of hyperinflation become intolerable. Over time, money loses its role as a store of value, unit of account, and medium of exchange. Barter becomes more common. And more stable unofficial monies—cigarettes or the U.S. dollar—start to replace the official money.

CASE STUDY

Life During the Bolivian Hyperinflation

The following article from the *Wall Street Journal* shows what life was like during the Bolivian hyperinflation of 1985.¹² What costs of inflation does this article emphasize?

Precarious Peso—Amid Wild Inflation, Bolivians Concentrate on Swapping Currency

LA PAZ, Bolivia—When Edgar Miranda gets his monthly teacher's pay of 25 million pesos, he hasn't a moment to lose. Every hour, pesos drop in value. So, while his wife rushes to market to lay in a month's supply of rice and noodles, he is off with the rest of the pesos to change them into black-market dollars.

Mr. Miranda is practicing the First Rule of Survival amid the most out-of-control inflation in the world today. Bolivia is a case study of how runaway inflation undermines a society. Price increases are so huge that the figures build up almost beyond comprehension. In one six-month period, for example, prices soared at an annual rate of 38,000%. By official count, however, last year's inflation reached 2,000%, and this year's is expected to hit 8,000%—though other estimates range many times higher. In any event, Bolivia's rate dwarfs Israel's 370% and Argentina's 1,100%—two other cases of severe inflation.

¹² Reprinted by permission of the *Wall Street Journal*, © August 13, 1985, page 1, Dow Jones & Company, Inc. All rights reserved worldwide.

It is easier to comprehend what happens to the 38-year-old Mr. Miranda's pay if he doesn't quickly change it into dollars. The day he was paid 25 million pesos, a dollar cost 500,000 pesos. So he received \$50. Just days later, with the rate at 900,000 pesos, he would have received \$27.

"We think only about today and converting every peso into dollars," says Ronald MacLean, the manager of a gold-mining firm. "We have become myopic."

And intent on survival. Civil servants won't hand out a form without a bribe. Lawyers, accountants, hairdressers, even prostitutes have almost given up working to become money-changers in the streets. Workers stage repeated strikes and steal from their bosses. The bosses smuggle production abroad, take out phony loans, duck taxes—anything to get dollars for speculation.

The production at the state mines, for example, dropped to 12,000 tons last year from 18,000. The miners pad their wages by smuggling out the richest ore in their lunch pails, and the ore goes by a contraband network into neighboring Peru. Without a major tin mine, Peru now exports some 4,000 metric tons of tin a year.


"We don't produce anything. We are all currency speculators," a heavy-equipment dealer in La Paz says. "People don't know what's good and bad anymore. We have become an amoral society. . . ."

It is an open secret that practically all of the black-market dollars come from the illegal cocaine trade with the U.S. Cocaine traffickers earn an estimated \$1 billion a year. . . .

But meanwhile the country is suffering from inflation largely because the government's revenues cover a mere 15% of its expenditures and its deficit has widened to nearly 25% of the country's total annual output. The revenues are hurt by a lag in tax payments, and taxes aren't being collected largely because of widespread theft and bribery.

The Causes of Hyperinflation

Why do hyperinflations start, and how do they end? This question can be answered at different levels.



© The New Yorker Collection 1997 Robert Mankoff from cartoonbank.com. All Rights Reserved.

"I told you the Fed should have tightened."

The most obvious answer is that hyperinflations are caused by excessive growth in the supply of money. When the central bank prints money, the price level rises. When it prints money rapidly enough, the result is hyperinflation. To stop the hyperinflation, the central bank must reduce the rate of money growth.

This answer is incomplete, however, for it leaves open the question of why central banks in hyperinflating economies choose to print so much money. To address this deeper question, we must turn our attention from monetary to fiscal policy. Most hyperinflations begin when the government has inadequate tax

revenue to pay for its spending. Although the government might prefer to finance this budget deficit by issuing debt, it may find itself unable to borrow, perhaps because lenders view the government as a bad credit risk. To cover the deficit, the government turns to the only mechanism at its disposal—the printing press. The result is rapid money growth and hyperinflation.

Once the hyperinflation is under way, the fiscal problems become even more severe. Because of the delay in collecting tax payments, real tax revenue falls as inflation rises. Thus, the government's need to rely on seigniorage is self-reinforcing. Rapid money creation leads to hyperinflation, which leads to a larger budget deficit, which leads to even more rapid money creation.

The ends of hyperinflations almost always coincide with fiscal reforms. Once the magnitude of the problem becomes apparent, the government musters the political will to reduce government spending and increase taxes. These fiscal reforms reduce the need for seigniorage, which allows a reduction in money growth. Hence, even if inflation is always and everywhere a monetary phenomenon, the end of hyperinflation is often a fiscal phenomenon as well.¹³

CASE STUDY

Hyperinflation in Interwar Germany

After World War I, Germany experienced one of history's most spectacular examples of hyperinflation. At the war's end, the Allies demanded that Germany pay substantial reparations. These payments led to fiscal deficits in Germany, which the German government eventually financed by printing large quantities of money.


Panel (a) of Figure 4-6 shows the quantity of money and the general price level in Germany from January 1922 to December 1924. During this period both money and prices rose at an amazing rate. For example, the price of a daily newspaper rose from 0.30 marks in January 1921 to 1 mark in May 1922, to 8 marks in October 1922, to 100 marks in February 1923, and to 1,000 marks in September 1923. Then, in the fall of 1923, prices took off: the newspaper sold for 2,000 marks on October 1, 20,000 marks on October 15, 1 million marks on October 29, 15 million marks on November 9, and 70 million marks on November 17. In December 1923 the money supply and prices abruptly stabilized.¹⁴

Just as fiscal problems caused the German hyperinflation, a fiscal reform ended it. At the end of 1923, the number of government employees was cut by one-third, and the reparations payments were temporarily suspended and eventually

¹³ For more on these issues, see Thomas J. Sargent, "The End of Four Big Inflations," in Robert Hall, ed., *Inflation* (Chicago: University of Chicago Press, 1983), 41–98; and Rudiger Dornbusch and Stanley Fischer, "Stopping Hyperinflations: Past and Present," *Weltwirtschaftliches Archiv* 122 (April 1986): 1–47.

¹⁴ The data on newspaper prices are from Michael Mussa, "Sticky Individual Prices and the Dynamics of the General Price Level," *Carnegie-Rochester Conference on Public Policy* 15 (Autumn 1981): 261–296.

figure 4-6



Money and Prices in Interwar Germany Panel (a) shows the money supply and the price level in Germany from January 1922 to December 1924. The immense increases in the money supply and the price level provide a dramatic illustration of the effects of printing large amounts of money. Panel (b) shows inflation and real money balances. As inflation rose, real money balances fell. When the inflation ended at the end of 1923, real money balances rose.

Source: Adapted from Thomas J. Sargent, "The End of Four Big Inflation," in Robert Hall, ed., *Inflation* (Chicago: University of Chicago Press, 1983): 41–98.

reduced. At the same time, a new central bank, the Rentenbank, replaced the old central bank, the Reichsbank. The Rentenbank was committed to not financing the government by printing money.

According to our theoretical analysis of money demand, an end to a hyperinflation should lead to an increase in real money balances as the cost of holding money falls. Panel (b) of Figure 4-6 shows that real money balances in Germany did fall as inflation increased, and then increased again as inflation fell. Yet the increase in real money balances was not immediate. Perhaps the adjustment of real money balances to the cost of holding money is a gradual process. Or perhaps it took time for people in Germany to believe that the inflation had ended, so that expected inflation fell more gradually than actual inflation.

4-8 Conclusion: The Classical Dichotomy

We have finished our discussion of money and inflation. Let's now step back and examine a key assumption that has been implicit in our discussion.

In Chapter 3, we explained many macroeconomic variables. Some of these variables were *quantities*, such as real GDP and the capital stock; others were *relative prices*, such as the real wage and the real interest rate. But all of these variables had one thing in common—they measured a physical (rather than a monetary) quantity. Real GDP is the quantity of goods and services produced in a given year, and the capital stock is the quantity of machines and structures available at a given time. The real wage is the quantity of output a worker earns for each hour of work, and the real interest rate is the quantity of output a person earns in the future by lending one unit of output today. All variables measured in physical units, such as quantities and relative prices, are called **real variables**.

In this chapter we examined **nominal variables**—variables expressed in terms of money. The economy has many nominal variables, such as the price level, the inflation rate, and the dollar wage a person earns.

At first it may seem surprising that we were able to explain real variables without introducing nominal variables or the existence of money. In Chapter 3 we studied the level and allocation of the economy's output without mentioning the price level or the rate of inflation. Our theory of the labor market explained the real wage without explaining the nominal wage.

Economists call this theoretical separation of real and nominal variables the **classical dichotomy**. It is the hallmark of classical macroeconomic theory. The classical dichotomy is an important insight, because it simplifies economic theory. In particular, it allows us to examine real variables, as we have done, while ignoring nominal variables. The classical dichotomy arises because, in classical economic theory, changes in the money supply do not influence real variables. This irrelevance of money for real variables is called **monetary neutrality**. For many purposes—in particular for studying long-run issues—monetary neutrality is approximately correct.

Yet monetary neutrality does not fully describe the world in which we live. Beginning in Chapter 9, we discuss departures from the classical model and monetary neutrality. These departures are crucial for understanding many macroeconomic phenomena, such as short-run economic fluctuations.

Summary

1. Money is the stock of assets used for transactions. It serves as a store of value, a unit of account, and a medium of exchange. Different sorts of assets are used as money: commodity money systems use an asset with intrinsic value, whereas fiat money systems use an asset whose sole function is to serve as money. In modern economies, a central bank such as the Federal Reserve is responsible for controlling the supply of money.
2. The quantity theory of money assumes that the velocity of money is stable and concludes that nominal GDP is proportional to the stock of money. Because the factors of production and the production function determine real GDP, the quantity theory implies that the price level is proportional to the quantity of money. Therefore, the rate of growth in the quantity of money determines the inflation rate.
3. Seigniorage is the revenue that the government raises by printing money. It is a tax on money holding. Although seigniorage is quantitatively small in most economies, it is often a major source of government revenue in economies experiencing hyperinflation.
4. The nominal interest rate is the sum of the real interest rate and the inflation rate. The Fisher effect says that the nominal interest rate moves one-for-one with expected inflation.
5. The nominal interest rate is the opportunity cost of holding money. Thus, one might expect the demand for money to depend on the nominal interest rate. If it does, then the price level depends on both the current quantity of money and the quantities of money expected in the future.
6. The costs of expected inflation include shoeleather costs, menu costs, the cost of relative price variability, tax distortions, and the inconvenience of making inflation corrections. In addition, unexpected inflation causes arbitrary redistributions of wealth between debtors and creditors. One possible benefit of inflation is that it improves the functioning of labor markets by allowing real wages to reach equilibrium levels without cuts in nominal wages.
7. During hyperinflations, most of the costs of inflation become severe. Hyperinflations typically begin when governments finance large budget deficits by printing money. They end when fiscal reforms eliminate the need for seigniorage.
8. According to classical economic theory, money is neutral: the money supply does not affect real variables. Therefore, classical theory allows us to study

how real variables are determined without any reference to the money supply. The equilibrium in the money market then determines the price level and, as a result, all other nominal variables. This theoretical separation of real and nominal variables is called the classical dichotomy.

KEY CONCEPTS

Inflation	Central bank	Seigniorage
Hyperinflation	Federal Reserve	Nominal and real interest rates
Money	Open-market operations	Fisher equation and Fisher effect
Store of value	Currency	<i>Ex ante</i> and <i>ex post</i> real interest rates
Unit of account	Demand deposits	Shoeleather costs
Medium of exchange	Quantity equation	Menu costs
Fiat money	Transactions velocity of money	Real and nominal variables
Commodity money	Income velocity of money	Classical dichotomy
Gold standard	Real money balances	Monetary neutrality
Money supply	Money demand function	
Monetary policy	Quantity theory of money	

QUESTIONS FOR REVIEW

- Describe the functions of money.
- What is fiat money? What is commodity money?
- Who controls the money supply and how?
- Write the quantity equation and explain it.
- What does the assumption of constant velocity imply?
- Who pays the inflation tax?
- If inflation rises from 6 to 8 percent, what happens to real and nominal interest rates according to the Fisher effect?
- List all the costs of inflation you can think of, and rank them according to how important you think they are.
- Explain the roles of monetary and fiscal policy in causing and ending hyperinflations.
- Define the terms “real variable” and “nominal variable,” and give an example of each.

PROBLEMS AND APPLICATIONS

- What are the three functions of money? Which of the functions do the following items satisfy? Which do they not satisfy?
 - A credit card
 - A painting by Rembrandt
 - A subway token
- In the country of Wiknam, the velocity of money is constant. Real GDP grows by 5 percent per year, the money stock grows by 14 percent per year, and the nominal interest rate is 11 percent. What is the real interest rate?
- A newspaper article written by the Associated Press in 1994 reported that the U.S. economy

was experiencing a low rate of inflation. It said that “low inflation has a downside: 45 million recipients of Social Security and other benefits will see their checks go up by just 2.8 percent next year.”

- a. Why does inflation affect the increase in Social Security and other benefits?
 - b. Is this effect a cost of inflation as the article suggests? Why or why not?
4. Suppose you are advising a small country (such as Bermuda) on whether to print its own money or to use the money of its larger neighbor (such as the United States). What are the costs and benefits of a national money? Does the relative political stability of the two countries have any role in this decision?
 5. During World War II, both Germany and England had plans for a paper weapon: they each printed the other’s currency, with the intention of dropping large quantities by airplane. Why might this have been an effective weapon?
 6. Calvin Coolidge once said that “inflation is repudiation.” What might he have meant by this? Do you agree? Why or why not? Does it matter whether the inflation is expected or unexpected?
 7. Some economic historians have noted that during the period of the gold standard, gold discoveries were most likely to occur after a long deflation. (The discoveries of 1896 are an example.) Why might this be true?
 8. Suppose that consumption depends on the level of real money balances (on the grounds that real money balances are part of wealth). Show that if real money balances depend on the nominal interest rate, then an increase in the rate of money growth affects consumption, investment, and the real interest rate. Does the nominal interest rate adjust more than one-for-one or less than one-for-one to expected inflation?
 This deviation from the classical dichotomy and the Fisher effect is called the “Mundell–Tobin effect.” How might you decide whether the Mundell–Tobin effect is important in practice?
 9. Use the Internet to identify a country that has had high inflation over the past year and another country that has had low inflation. (*Hint:* One useful Web site is www.economist.com.) For these two countries, find the rate of money growth and the current level of the nominal interest rate. Relate your findings to the theories presented in this chapter.

The Cagan Model: How Current and Future Money Affect the Price Level

In this chapter we showed that if the quantity of real money balances demanded depends on the cost of holding money, the price level depends on both the current money supply and the future money supply. This appendix develops the *Cagan model* to show more explicitly how this works.¹⁵

To keep the math as simple as possible, we posit a money demand function that is linear in the natural logarithms of all the variables. The money demand function is

$$m_t - p_t = -\gamma(p_{t+1} - p_t), \quad (\text{A1})$$

where m_t is the log of the quantity of money at time t , p_t is the log of the price level at time t , and γ is a parameter that governs the sensitivity of money demand to the rate of inflation. By the property of logarithms, $m_t - p_t$ is the log of real money balances, and $p_{t+1} - p_t$ is the inflation rate between period t and period $t + 1$. This equation states that if inflation goes up by 1 percentage point, real money balances fall by γ percent.

We have made a number of assumptions in writing the money demand function in this way. First, by excluding the level of output as a determinant of money demand, we are implicitly assuming that it is constant. Second, by including the rate of inflation rather than the nominal interest rate, we are assuming that the real interest rate is constant. Third, by including actual inflation rather than expected inflation, we are assuming perfect foresight. All of these assumptions are to keep the analysis as simple as possible.

We want to solve Equation A1 to express the price level as a function of current and future money. To do this, note that Equation A1 can be rewritten as

$$p_t = \left(\frac{1}{1 + \gamma} \right) m_t + \left(\frac{\gamma}{1 + \gamma} \right) p_{t+1}. \quad (\text{A2})$$

This equation states that the current price level p_t is a weighted average of the current money supply m_t and the next period's price level p_{t+1} . The next period's price level will be determined the same way as this period's price level:

$$p_{t+1} = \left(\frac{1}{1 + \gamma} \right) m_{t+1} + \left(\frac{\gamma}{1 + \gamma} \right) p_{t+2}. \quad (\text{A3})$$

Now substitute Equation A3 for p_{t+1} in Equation A2 to obtain

$$p_t = \frac{1}{1 + \gamma} m_t + \frac{\gamma}{(1 + \gamma)^2} m_{t+1} + \frac{\gamma^2}{(1 + \gamma)^2} p_{t+2}. \quad (\text{A4})$$

¹⁵ This model is derived from Phillip Cagan, "The Monetary Dynamics of Hyperinflation," in Milton Friedman, ed., *Studies in the Quantity Theory of Money* (Chicago: University of Chicago Press, 1956).

Equation A4 states that the current price level is a weighted average of the current money supply, the next period's money supply, and the following period's price level. Once again, the price level in $t + 2$ is determined as in Equation A2:

$$p_{t+2} = \left(\frac{1}{1 + \gamma} \right) m_{t+2} + \left(\frac{\gamma}{1 + \gamma} \right) p_{t+3}. \quad (\text{A5})$$

Now substitute Equation A5 into Equation A4 to obtain

$$p_t = \frac{1}{1 + \gamma} m_t + \frac{\gamma}{(1 + \gamma)^2} m_{t+1} + \frac{\gamma^2}{(1 + \gamma)^3} m_{t+2} + \frac{\gamma^3}{(1 + \gamma)^3} p_{t+3}. \quad (\text{A6})$$

By now you see the pattern. We can continue to use Equation A2 to substitute for the future price level. If we do this an infinite number of times, we find

$$p_t = \left(\frac{1}{1 + \gamma} \right) \left[m_t + \left(\frac{\gamma}{1 + \gamma} \right) m_{t+1} + \left(\frac{\gamma}{1 + \gamma} \right)^2 m_{t+2} + \left(\frac{\gamma}{1 + \gamma} \right)^3 m_{t+3} + \dots \right], \quad (\text{A7})$$

where $\dots$ indicates an infinite number of analogous terms. According to Equation A7, the current price level is a weighted average of the current money supply and all future money supplies.

Note the importance of γ , the parameter governing the sensitivity of real money balances to inflation. The weights on the future money supplies decline geometrically at rate $\gamma/(1 + \gamma)$. If γ is small, then $\gamma/(1 + \gamma)$ is small, and the weights decline quickly. In this case, the current money supply is the primary determinant of the price level. (Indeed, if γ equals zero, then we obtain the quantity theory of money: the price level is proportional to the current money supply, and the future money supplies do not matter at all.) If γ is large, then $\gamma/(1 + \gamma)$ is close to 1, and the weights decline slowly. In this case, the future money supplies play a key role in determining today's price level.

Finally, let's relax the assumption of perfect foresight. If the future is not known with certainty, then we should write the money demand function as

$$m_t - p_t = -\gamma(Ep_{t+1} - p_t), \quad (\text{A8})$$

where Ep_{t+1} is the expected price level. Equation A8 states that real money balances depend on expected inflation. By following steps similar to those preceding, we can show that

$$p_t = \left(\frac{1}{1 + \gamma} \right) \left[m_t + \left(\frac{\gamma}{1 + \gamma} \right) Em_{t+1} + \left(\frac{\gamma}{1 + \gamma} \right)^2 Em_{t+2} + \left(\frac{\gamma}{1 + \gamma} \right)^3 Em_{t+3} + \dots \right]. \quad (\text{A9})$$

Equation A9 states that the price level depends on the current money supply and expected future money supplies.

Some economists use this model to argue that *credibility* is important for ending hyperinflation. Because the price level depends on both current and expected future money, inflation depends on both current and expected future money growth. Therefore, to end high inflation, both money growth and expected money growth must fall. Expectations, in turn, depend on credibility—the perception that the central bank is committed to a new, more stable policy.

How can a central bank achieve credibility in the midst of hyperinflation? Credibility is often achieved by removing the underlying cause of the hyperinflation—the need for seigniorage. Thus, a credible fiscal reform is often necessary for a credible change in monetary policy. This fiscal reform might take the form of reducing government spending and making the central bank more independent from the government. Reduced spending decreases the need for seigniorage, while increased independence allows the central bank to resist government demands for seigniorage.

MORE PROBLEMS AND APPLICATIONS

1. In the Cagan model, if the money supply is expected to grow at some constant rate μ (so that $Em_{t+s} = m_t + s\mu$), then Equation A9 can be shown to imply that $p_t = m_t + \gamma\mu$.
 - a. Interpret this result.
 - b. What happens to the price level p_t when the money supply m_t changes, holding the money growth rate μ constant?
 - c. What happens to the price level p_t when the money growth rate μ changes, holding the current money supply m_t constant?
 - d. If a central bank is about to reduce the rate of money growth μ but wants to hold the price level p_t constant, what should it do with m_t ? Can you see any practical problems that might arise in following such a policy?
 - e. How do your previous answers change in the special case where money demand does not depend on the expected rate of inflation (so that $\gamma = 0$)?

The Open Economy

No nation was ever ruined by trade.

— Benjamin Franklin


Even if you never leave your home town, you are an active participant in a global economy. When you go to the grocery store, for instance, you might choose between apples grown locally and grapes grown in Chile. When you make a deposit into your local bank, the bank might lend those funds to your next-door neighbor or to a Japanese company building a factory outside Tokyo. Because our economy is integrated with many others around the world, consumers have more goods and services from which to choose, and savers have more opportunities to invest their wealth.

In previous chapters we simplified our analysis by assuming a closed economy. In actuality, however, most economies are open: they export goods and services abroad, they import goods and services from abroad, and they borrow and lend in world financial markets. Figure 5-1 gives some sense of the importance of these international interactions by showing imports and exports as a percentage of GDP for seven major industrial countries. As the figure shows, imports and exports in the United States are more than 10 percent of GDP. Trade is even more important for many other countries—in Canada and the United Kingdom, for instance, imports and exports are about a third of GDP. In these countries, international trade is central to analyzing economic developments and formulating economic policies.

This chapter begins our study of open-economy macroeconomics. We begin in Section 5-1 with questions of measurement. To understand how the open economy works, we must understand the key macroeconomic variables that measure the interactions among countries. Accounting identities reveal a key insight: the flow of goods and services across national borders is always matched by an equivalent flow of funds to finance capital accumulation.

In Section 5-2 we examine the determinants of these international flows. We develop a model of the small open economy that corresponds to our model of the closed economy in Chapter 3. The model shows the factors that determine whether a country is a borrower or a lender in world markets, and how policies at home and abroad affect the flows of capital and goods.

figure 5-1



Imports and Exports as a Percentage of Output: 2000 While international trade is important for the United States, it is even more vital for other countries.

Source: OECD.

In Section 5-3 we extend the model to discuss the prices at which a country makes exchanges in world markets. We examine what determines the price of domestic goods relative to foreign goods. We also examine what determines the rate at which the domestic currency trades for foreign currencies. Our model shows how protectionist trade policies—policies designed to protect domestic industries from foreign competition—influence the amount of international trade and the exchange rate.

5-1 The International Flows of Capital and Goods

The key macroeconomic difference between open and closed economies is that, in an open economy, a country's spending in any given year need not equal its output of goods and services. A country can spend more than it produces by borrowing from abroad, or it can spend less than it produces and lend the difference to foreigners. To understand this more fully, let's take another look at national income accounting, which we first discussed in Chapter 2.

The Role of Net Exports

Consider the expenditure on an economy's output of goods and services. In a closed economy, all output is sold domestically, and expenditure is divided into three components: consumption, investment, and government purchases. In an open economy, some output is sold domestically and some is exported to be sold abroad. We can divide expenditure on an open economy's output Y into four components:

- C^d , consumption of domestic goods and services,
- I^d , investment in domestic goods and services,
- G^d , government purchases of domestic goods and services,
- EX , exports of domestic goods and services.

The division of expenditure into these components is expressed in the identity

$$Y = C^d + I^d + G^d + EX.$$

The sum of the first three terms, $C^d + I^d + G^d$, is domestic spending on domestic goods and services. The fourth term, EX , is foreign spending on domestic goods and services.

We now want to make this identity more useful. To do this, note that domestic spending on all goods and services is the sum of domestic spending on domestic goods and services and on foreign goods and services. Hence, total consumption C equals consumption of domestic goods and services C^d plus consumption of foreign goods and services C^f ; total investment I equals investment in domestic goods and services I^d plus investment in foreign goods and services I^f ; and total government purchases G equals government purchases of domestic goods and services G^d plus government purchases of foreign goods and services G^f . Thus,

$$\begin{aligned} C &= C^d + C^f, \\ I &= I^d + I^f, \\ G &= G^d + G^f. \end{aligned}$$

We substitute these three equations into the identity above:

$$Y = (C - C^f) + (I - I^f) + (G - G^f) + EX.$$

We can rearrange to obtain

$$Y = C + I + G + EX - (C^f + I^f + G^f).$$

The sum of domestic spending on foreign goods and services ($C^f + I^f + G^f$) is expenditure on imports (IM). We can thus write the national income accounts identity as

$$Y = C + I + G + EX - IM.$$

Because spending on imports is included in domestic spending ($C + I + G$), and because goods and services imported from abroad are not part of a country's

output, this equation subtracts spending on imports. Defining **net exports** to be exports minus imports ($NX = EX - IM$), the identity becomes

$$Y = C + I + G + NX.$$

This equation states that expenditure on domestic output is the sum of consumption, investment, government purchases, and net exports. This is the most common form of the national income accounts identity; it should be familiar from Chapter 2.

The national income accounts identity shows how domestic output, domestic spending, and net exports are related. In particular,

$$NX = Y - (C + I + G)$$

Net Exports = Output – Domestic Spending.

This equation shows that in an open economy, domestic spending need not equal the output of goods and services. *If output exceeds domestic spending, we export the difference: net exports are positive. If output falls short of domestic spending, we import the difference: net exports are negative.*

International Capital Flows and the Trade Balance

In an open economy, as in the closed economy we discussed in Chapter 3, financial markets and goods markets are closely related. To see the relationship, we must rewrite the national income accounts identity in terms of saving and investment. Begin with the identity

$$Y = C + I + G + NX.$$

Subtract C and G from both sides to obtain

$$Y - C - G = I + NX.$$

Recall from Chapter 3 that $Y - C - G$ is national saving S , the sum of private saving, $Y - T - C$, and public saving, $T - G$. Therefore,

$$S = I + NX.$$

Subtracting I from both sides of the equation, we can write the national income accounts identity as

$$S - I = NX.$$

This form of the national income accounts identity shows that an economy's net exports must always equal the difference between its saving and its investment.

Let's look more closely at each part of this identity. The easy part is the right-hand side, NX , which is our net export of goods and services. Another name for net exports is the **trade balance**, because it tells us how our trade in goods and services departs from the benchmark of equal imports and exports.

The left-hand side of the identity is the difference between domestic saving and domestic investment, $S - I$, which we'll call **net capital outflow**. (It's sometimes called *net foreign investment*.) If net capital outflow is positive, our saving exceeds our investment, and we are lending the excess to foreigners. If the net capital outflow is negative, our investment exceeds our saving, and we are financing this extra investment by borrowing from abroad. Thus, net capital outflow equals the amount that domestic residents are lending abroad minus the amount that foreigners are lending to us. It reflects the international flow of funds to finance capital accumulation.

The national income accounts identity shows that net capital outflow always equals the trade balance. That is,

$$\begin{array}{rcl} \text{Net Capital Outflow} & = & \text{Trade Balance} \\ S - I & = & NX. \end{array}$$

If $S - I$ and NX are positive, we have a **trade surplus**. In this case, we are net lenders in world financial markets, and we are exporting more goods than we are importing. If $S - I$ and NX are negative, we have a **trade deficit**. In this case, we are net borrowers in world financial markets, and we are importing more goods than we are exporting. If $S - I$ and NX are exactly zero, we are said to have **balanced trade** because the value of imports equals the value of exports.

The national income accounts identity shows that the international flow of funds to finance capital accumulation and the international flow of goods and services are two sides of the same coin. On the one hand, if our saving exceeds our investment, the saving that is not invested domestically is used to make loans to foreigners. Foreigners require these loans because we are providing them with more goods and services than they are providing us. That is, we are running a trade surplus. On the other hand, if our investment exceeds our saving, the extra investment must be financed by borrowing from abroad. These foreign loans enable us to import more

table 5-1

International Flows of Goods and Capital: Summary

This table shows the three outcomes that an open economy can experience.

Trade Surplus	Balanced Trade	Trade Deficit
Exports > Imports	Exports = Imports	Exports < Imports
Net Exports > 0	Net Exports = 0	Net Exports < 0
$Y > C + I + G$	$Y = C + I + G$	$Y < C + I + G$
Savings > Investment	Saving = Investment	Saving < Investment
Net Capital Outflow > 0	Net Capital Outflow = 0	Net Capital Outflow < 0



International Flows of Goods and Capital: An Example

The equality of net exports and net capital outflow is an identity: it must hold by the way the numbers are added up. But it is easy to miss the intuition behind this important relationship. The best way to understand it is to consider an example.

Imagine that Bill Gates sells a copy of the Windows operating system to a Japanese consumer for 5,000 yen. Because Mr. Gates is a U.S. resident, the sale represents an export of the United States. Other things equal, U.S. net exports rise. What else happens to make the identity hold? It depends on what Mr. Gates does with the 5,000 yen.

Suppose Mr. Gates decides to stuff the 5,000 yen in his mattress. In this case, Mr. Gates has allocated some of his saving to an investment in the Japanese economy (in the form of the Japanese currency) rather than to an investment in the U.S. economy. Thus, U.S. saving exceeds U.S. investment. The rise in U.S. net exports is matched by a rise in the U.S. net capital outflow.

If Mr. Gates wants to invest in Japan, however, he is unlikely to make currency his asset of choice. He might use the 5,000 yen to buy some stock in, say, the Sony Corporation, or he might buy a bond issued by the Japanese government. In either case, some of U.S. saving is flowing abroad. Once again, the U.S. net capital outflow exactly balances U.S. net exports.

The opposite situation occurs in Japan. When the Japanese consumer buys a copy of the Windows operating system, Japan's purchases of goods and services ($C + I + G$) rise, but there is no change in what Japan has produced (Y). The transaction reduces Japan's saving ($S = Y - C - G$) for a given level of investment (I). While the U.S. experiences a net capital outflow, Japan experiences a net capital inflow.

Now let's change the example. Suppose that instead of investing his 5,000 yen in a Japanese asset, Mr. Gates uses it to buy something made in Japan, such as a Sony Walkman. In this case, imports into the United States rise. Together, the Windows export and the Walkman import represent balanced trade between Japan and the United States. Because exports and imports rise equally, net exports and net capital outflow are both unchanged.

A final possibility is that Mr. Gates exchanges his 5,000 yen for U.S. dollars at a local bank. But this doesn't change the situation: the bank now has to do something with the 5,000 yen. It can buy Japanese assets (a U.S. net capital outflow); it can buy a Japanese good (a U.S. import); or it can sell the yen to another American who wants to make such a transaction. If you follow the money, you can see that, in the end, U.S. net exports must equal U.S. net capital outflow.

goods and services than we export. That is, we are running a trade deficit. Table 5-1 summarizes these lessons.

Note that the international flow of capital can take many forms. It is easiest to assume—as we have done so far—that when we run a trade deficit, foreigners make loans to us. This happens, for example, when the Japanese buy the debt issued by U.S. corporations or by the U.S. government. But the flow of capital can also take the form of foreigners buying domestic assets, such as when a citizen of Germany buys stock from an American on the New York Stock Exchange. Whether foreigners are buying domestically issued debt or domestically owned assets, they are obtaining a claim to the future returns to domestic capital. In both cases, foreigners end up owning some of the domestic capital stock.

5-2 Saving and Investment in a Small Open Economy

So far in our discussion of the international flows of goods and capital, we have merely rearranged accounting identities. That is, we have defined some of the variables that measure transactions in an open economy, and we have shown the links among these variables that follow from their definitions. Our next step is to develop a model that explains the behavior of these variables. We can then use the model to answer questions such as how the trade balance responds to changes in policy.

Capital Mobility and the World Interest Rate

In a moment we present a model of the international flows of capital and goods. Because the trade balance equals the net capital outflow, which in turn equals saving minus investment, our model focuses on saving and investment. To develop this model, we use some elements that should be familiar from Chapter 3, but in contrast to the Chapter 3 model, we do not assume that the real interest rate equilibrates saving and investment. Instead, we allow the economy to run a trade deficit and borrow from other countries, or to run a trade surplus and lend to other countries.

If the real interest rate does not adjust to equilibrate saving and investment in this model, what *does* determine the real interest rate? We answer this question here by considering the simple case of a **small open economy** with perfect capital mobility. By “small” we mean that this economy is a small part of the world market and thus, by itself, can have only a negligible effect on the world interest rate. By “perfect capital mobility” we mean that residents of the country have full access to world financial markets. In particular, the government does not impede international borrowing or lending.

Because of this assumption of perfect capital mobility, the interest rate in our small open economy, r , must equal the **world interest rate** r^* , the real interest rate prevailing in world financial markets:

$$r = r^*.$$

Residents of the small open economy need never borrow at any interest rate above r^* , because they can always get a loan at r^* from abroad. Similarly, residents of this economy need never lend at any interest rate below r^* because they can always earn r^* by lending abroad. Thus, the world interest rate determines the interest rate in our small open economy.

Let us discuss for a moment what determines the world real interest rate. In a closed economy, the equilibrium of domestic saving and domestic investment determines the interest rate. Barring interplanetary trade, the world economy is a closed economy. Therefore, the equilibrium of world saving and world investment determines the world interest rate. Our small open economy has a negligible effect on the world real interest rate because, being a small part of the world,

it has a negligible effect on world saving and world investment. Hence, our small open economy takes the world interest rate as exogenously given.

The Model

To build the model of the small open economy, we take three assumptions from Chapter 3:

- The economy's output Y is fixed by the factors of production and the production function. We write this as

$$Y = \bar{Y} = F(\bar{K}, \bar{L}).$$

- Consumption C is positively related to disposable income $Y - T$. We write the consumption function as

$$C = C(Y - T).$$

- Investment I is negatively related to the real interest rate r . We write the investment function as

$$I = I(r).$$

These are the three key parts of our model. If you do not understand these relationships, review Chapter 3 before continuing.

We can now return to the accounting identity and write it as

$$\begin{aligned} NX &= (Y - C - G) - I \\ NX &= S - I. \end{aligned}$$

Substituting our three assumptions from Chapter 3 and the condition that the interest rate equals the world interest rate, we obtain


$$\begin{aligned} NX &= [\bar{Y} - C(\bar{Y} - T) - G] - I(r^*) \\ &= \bar{S} - I(r^*). \end{aligned}$$

This equation shows what determines saving S and investment I —and thus the trade balance NX . Remember that saving depends on fiscal policy: lower government purchases G or higher taxes T raise national saving. Investment depends on the world real interest rate r^* ; high interest rates make some investment projects unprofitable. Therefore, the trade balance depends on these variables as well.

In Chapter 3 we graphed saving and investment as in Figure 5-2. In the closed economy studied in that chapter, the real interest rate adjusts to equilibrate saving and investment—that is, the real interest rate is found where the saving and investment curves cross. In the small open economy, however, the real interest rate equals the world real interest rate. *The trade balance is determined by the difference between saving and investment at the world interest rate.*

At this point, you might wonder about the mechanism that causes the trade balance to equal the net capital outflow. The determinants of the capital flows are

figure 5-2



Saving and Investment in a Small Open Economy In a closed economy, the real interest rate adjusts to equilibrate saving and investment. In a small open economy, the interest rate is determined in world financial markets. The difference between saving and investment determines the trade balance. Here there is a trade surplus, because at the world interest rate, saving exceeds investment.

easy to understand. When saving falls short of investment, investors borrow from abroad; when saving exceeds investment, the excess is lent to other countries. But what causes those who import and export to behave in a way that ensures that the international flow of goods exactly balances this international flow of capital? For now we leave this question unanswered, but we return to it in Section 5-3 when we discuss the determination of exchange rates.

How Policies Influence the Trade Balance


Suppose that the economy begins in a position of balanced trade. That is, at the world interest rate, investment I equals saving S , and net exports NX equal zero. Let's use our model to predict the effects of government policies at home and abroad.

Fiscal Policy at Home Consider first what happens to the small open economy if the government expands domestic spending by increasing government purchases. The increase in G reduces national saving, because $S = Y - C - G$. With an unchanged world real interest rate, investment remains the same. Therefore, saving falls below investment, and some investment must now be financed by borrowing from abroad. Because $NX = S - I$, the fall in S implies a fall in NX . The economy now runs a trade deficit.

The same logic applies to a decrease in taxes. A tax cut lowers T , raises disposable income $Y - T$, stimulates consumption, and reduces national saving. (Even though some of the tax cut finds its way into private saving, public saving falls by the full amount of the tax cut; in total, saving falls.) Because $NX = S - I$, the reduction in national saving in turn lowers NX .

Figure 5-3 illustrates these effects. A fiscal-policy change that increases private consumption C or public consumption G reduces national saving ($Y - C - G$) and, therefore, shifts the vertical line that represents saving from S_1 to S_2 . Because NX is the distance between the saving schedule and the investment schedule at the world interest rate, this shift reduces NX . Hence, starting from balanced trade, a change in fiscal policy that reduces national saving leads to a trade deficit.

figure 5-3




A Fiscal Expansion at Home in a Small Open Economy An increase in government purchases or a reduction in taxes reduces national saving and thus shifts the saving schedule to the left, from S_1 to S_2 . The result is a trade deficit.

Fiscal Policy Abroad Consider now what happens to a small open economy when foreign governments increase their government purchases. If these foreign countries are a small part of the world economy, then their fiscal change has a negligible impact on other countries. But if these foreign countries are a large part of the world economy, their increase in government purchases reduces world saving and causes the world interest rate to rise.

The increase in the world interest rate raises the cost of borrowing and, thus, reduces investment in our small open economy. Because there has been no change in domestic saving, saving S now exceeds investment I , and some of our saving begins to flow abroad. Since $NX = S - I$, the reduction in I must also increase NX . Hence, reduced saving abroad leads to a trade surplus at home.

Figure 5-4 illustrates how a small open economy starting from balanced trade responds to a foreign fiscal expansion. Because the policy change is occurring

figure 5-4




A Fiscal Expansion Abroad in a Small Open Economy A fiscal expansion in a foreign economy large enough to influence world saving and investment raises the world interest rate from r_1^* to r_2^* . The higher world interest rate reduces investment in this small open economy, causing a trade surplus.

abroad, the domestic saving and investment schedules remain the same. The only change is an increase in the world interest rate from r_1^* to r_2^* . The trade balance is the difference between the saving and investment schedules; because saving exceeds investment at r_2^* , there is a trade surplus. Hence, an increase in the world interest rate due to a fiscal expansion abroad leads to a trade surplus.

Shifts in Investment Demand Consider what happens to our small open economy if its investment schedule shifts outward—that is, if the demand for investment goods at every interest rate increases. This shift would occur if, for example, the government changed the tax laws to encourage investment by providing an investment tax credit. Figure 5-5 illustrates the impact of a shift in the investment schedule. At a given world interest rate, investment is now higher. Because saving is unchanged, some investment must now be financed by borrowing from abroad, which means the net capital outflow is negative. Put differently, because $NX = S - I$, the increase in I implies a decrease in NX . Hence, an outward shift in the investment schedule causes a trade deficit.

figure 5-5



A Shift in the Investment Schedule in a Small Open Economy An outward shift in the investment schedule from $I(r)_1$ to $I(r)_2$ increases the amount of investment at the world interest rate r^* . As a result, investment now exceeds saving, which means the economy is borrowing from abroad and running a trade deficit.

Evaluating Economic Policy

Our model of the open economy shows that the flow of goods and services measured by the trade balance is inextricably connected to the international flow of funds for capital accumulation. The net capital outflow is the difference between domestic saving and domestic investment. Thus, the impact of economic policies on the trade balance can always be found by examining their impact on domestic saving and domestic investment. Policies that increase investment or decrease saving tend to cause a trade deficit, and policies that decrease investment or increase saving tend to cause a trade surplus.

Our analysis of the open economy has been positive, not normative. That is, our analysis of how economic policies influence the international flows of capital and goods has not told us whether these policies are desirable. Evaluating economic policies and their impact on the open economy is a frequent topic of debate among economists and policymakers.

When a country runs a trade deficit, policymakers must confront the question of whether it represents a national problem. Most economists view a trade deficit not as a problem in itself, but perhaps as a symptom of a problem. A trade deficit could be a reflection of low saving. In a closed economy, low saving leads to low investment and a smaller future capital stock. In an open economy, low saving leads to a trade deficit and a growing foreign debt, which eventually must be repaid. In both cases, high current consumption leads to lower future consumption, implying that future generations bear the burden of low national saving.

Yet trade deficits are not always a reflection of economic malady. When poor rural economies develop into modern industrial economies, they sometimes finance their high levels of investment with foreign borrowing. In these cases, trade deficits are a sign of economic development. For example, South Korea ran large trade deficits throughout the 1970s, and it became one of the success stories of economic growth. The lesson is that one cannot judge economic performance from the trade balance alone. Instead, one must look at the underlying causes of the international flows.

CASE STUDY


The U.S. Trade Deficit

During the 1980s and 1990s, the United States ran large trade deficits. Panel (a) of Figure 5-6 documents this experience by showing net exports as a percentage of GDP. The exact size of the trade deficit fluctuated over time, but it was large throughout these two decades. In 2000, the trade deficit was \$371 billion, or 3.7 percent of GDP. As accounting identities require, this trade deficit had to be financed by borrowing from abroad (or, equivalently, by selling U.S. assets abroad). During this period, the United States went from being the world's largest creditor to the world's largest debtor.

What caused the U.S. trade deficit? There is no single explanation. But to understand some of the forces at work, it helps to look at national saving and domestic investment, as shown in panel (b) of the figure. Keep in mind that the trade deficit is the difference between saving and investment.

The start of the trade deficit coincided with a fall in national saving. This development can be explained by the expansionary fiscal policy in the 1980s. With the support of President Reagan, the U.S. Congress passed legislation in 1981 that substantially cut personal income taxes over the next three years. Because these tax cuts were not met with equal cuts in government spending, the federal budget went into deficit. These budget deficits were among the largest ever experienced in a period of peace and prosperity, and they continued long after Reagan left office. According to our model, such a policy should reduce national

figure 5-6



The Trade Balance, Saving, and Investment: The U.S. Experience Panel (a) shows the trade balance as a percentage of GDP. Positive numbers represent a surplus, and negative numbers represent a deficit. Panel (b) shows national saving and investment as a percentage of GDP since 1960. The trade balance equals saving minus investment.

Source: U.S. Department of Commerce.

saving, thereby causing a trade deficit. And, in fact, that is exactly what happened. Because the government budget and trade balance went into deficit at roughly the same time, these shortfalls were called the *twin deficits*.

Things started to change in the 1990s, when the U.S. federal government got its fiscal house in order. The first President Bush and President Clinton both signed tax increases, while Congress kept a lid on spending. In addition to these policy changes, rapid productivity growth in the late 1990s raised incomes and, thus, further increased tax revenue. These developments moved the U.S. federal budget from deficit to surplus, which in turn caused national saving to rise.

In contrast to what our model predicts, the increase in national saving did not coincide with a shrinking trade deficit, because domestic investment rose at the same time. The likely explanation is that the boom in information technology caused an expansionary shift in the U.S. investment function. Even though fiscal policy was pushing the trade deficit toward surplus, the investment boom was an even stronger force pushing the trade balance toward deficit.

The history of the U.S. trade deficit shows that this statistic, by itself, does not tell us much about what is happening in the economy. We have to look deeper at saving, investment, and the policies and events that cause them to change over time.

5-3 Exchange Rates

Having examined the international flows of capital and of goods and services, we now extend the analysis by considering the prices that apply to these transactions. The *exchange rate* between two countries is the price at which residents of those countries trade with each other. In this section we first examine precisely what the exchange rate measures, and we then discuss how exchange rates are determined.

Nominal and Real Exchange Rates

Economists distinguish between two exchange rates: the nominal exchange rate and the real exchange rate. Let's discuss each in turn and see how they are related.

The Nominal Exchange Rate The **nominal exchange rate** is the relative price of the currency of two countries. For example, if the exchange rate between the U.S. dollar and the Japanese yen is 120 yen per dollar, then you can exchange one dollar for 120 yen in world markets for foreign currency. A Japanese who wants to obtain dollars would pay 120 yen for each dollar he bought. An American who wants to obtain yen would get 120 yen for each dollar he paid. When people refer to “the exchange rate” between two countries, they usually mean the nominal exchange rate.



How Newspapers Report the Exchange Rate

You can find nominal exchange rates reported daily in many newspapers. Here's how they are reported in the *Wall Street Journal*.

Notice that each exchange rate is reported in two ways. On this Thursday, 1 dollar bought 116.29 yen, and 1 yen bought 0.008599 dollars. We can say the exchange rate is 116.29 yen per dollar, or we can say the exchange rate is 0.008599 dollars per yen. Because 0.008599 equals 1/116.29, these two ways of expressing

the exchange rate are equivalent. This book always expresses the exchange rate in units of foreign currency per dollar.

The exchange rate on this Thursday of 116.29 yen per dollar was down from 117.67 yen per dollar on Wednesday. Such a fall in the exchange rate is called a *depreciation* of the dollar; a rise in the exchange rate is called an *appreciation*. When the domestic currency depreciates, it buys less of the foreign currency; when it appreciates, it buys more.

CURRENCY TRADING

Thursday, September 20, 2001

EXCHANGE RATES

The foreign exchange mid-range rates below apply to trading among banks in amounts of \$1 million and more, as quoted at 4 p.m. Eastern time by Reuters and other sources. Retail transactions provide fewer units of foreign currency per dollar. Rates for the 12 Euro currency countries are derived from the latest dollar-euro rate using the exchange ratios set 1/1/99.

Country	U.S. \$ EQUIV.		CURRENCY PER U.S. \$	
	Thu	Wed	Thu	Wed
Argentina (Peso)	1.0002	1.0003	.9998	.9997
Australia (Dollar)	.4932	.4931	2.0274	2.0278
Austria (Schilling)	.06740	.06735	14.837	14.848
Bahrain (Dinar)	2.6525	2.6518	.3770	.3771
Belgium (Franc)	.0230	.0230	43.4955	43.5284
Brazil (Real)	.3612	.3699	2.7685	2.7035
Britain (Pound)	1.4671	1.4684	.6816	.6810
1-month forward	1.4648	1.4658	.6827	.6822
3-months forward	1.4603	1.4615	.6848	.6842
6-months forward	1.4540	1.4549	.6878	.6873
Canada (Dollar)	.6364	.6376	1.5713	1.5683
1-month forward	.6359	.6371	1.5725	1.5696
3-months forward	.6354	.6365	1.5739	1.5712
6-months forward	.6348	.6359	1.5752	1.5725
Chile (Peso)	.001440	.001449	694.55	690.15
China (Renminbi)	.1208	.1208	8.2766	8.2766
Colombia (Peso)	.0004269	.0004264	2342.50	2345.00
Czech Rep. (Koruna)				
Commercial rate	.02704	.02716	36.985	36.822
Denmark (Krone)	.1246	.1245	8.0260	8.0315
Ecuador (US Dollar)-e	1.0000	1.0000	1.0000	1.0000
Finland (Markka)	.1560	.1559	6.4108	6.4157
France (Franc)	.1414	.1413	7.0727	7.0780
1-month forward	.1413	.1412	7.0794	7.0846
3-months forward	.1410	.1409	7.0921	7.0970
6-months forward	.1407	.1406	7.1088	7.1146
Germany (Mark)	.4742	.4738	2.1088	2.1104
1-month forward	.4738	.4734	2.1108	2.1124
3-months forward	.4729	.4726	2.1146	2.1161
6-months forward	.4718	.4714	2.1196	2.1214
Greece (Drachma)	.002722	.002720	367.41	367.70
Hong Kong (Dollar)	.1282	.1282	7.7995	7.7992
Hungary (Forint)	.003589	.003592	278.67	278.40
India (Rupee)	.02084	.02083	47.980	48.010
Indonesia (Rupiah)	.0001059	.0001041	9445	9605
Ireland (Punt)	1.1776	1.1767	.8492	.8498
Israel (Shekel)	.2297	.2304	4.3541	4.3400
Italy (Lira)	.0004790	.0004786	2087.74	2089.31

Country	U.S. \$ EQUIV.		CURRENCY PER U.S. \$	
	Thu	Wed	Thu	Wed
Japan (Yen)	.008599	.008498	116.29	117.67
1-month forward	.008618	.008517	116.04	117.41
3-months forward	.008655	.008553	115.53	116.92
6-months forward	.008707	.008603	114.85	116.23
Jordan (Dinar)	1.4069	1.4069	.7108	.7108
Kuwait (Dinar)	3.2830	3.2841	.3046	.3045
Lebanon (Pound)	.0006604	.0006604	1514.25	1514.25
Malaysia (Ringgit)-b	.2632	.2632	3.8000	3.8000
Malta (Lira)	2.779	2.2774	.4390	.4391
Mexico (Peso)				
Floating rate	.1055	.1060	9.4825	9.4325
Netherlands (Guilder)	.4209	.4205	2.3761	2.3779
New Zealand (Dollar)	.4128	.4122	2.4225	2.4260
Norway (Krone)	.1169	.1169	8.5530	8.5551
Pakistan (Rupee)	.01559	.01559	64.150	64.150
Peru (new Sol)	.2857	.2857	3.5005	3.5003
Philippines (Peso)	.01949	.01946	51.300	51.375
Poland (Zloty)-d	.2404	.2382	4.1600	4.1975
Portugal (Escudo)	.004626	.004623	216.16	216.33
Russia (Ruble)-a	.03399	.03397	29.422	29.442
Saudi Arabia (Riyal)	.2666	.2666	3.7508	3.7505
Singapore (Dollar)	.5758	.5749	1.7368	1.7395
Slovak Rep. (Koruna)	.02116	.02118	47.253	47.219
South Africa (Rand)	.1149	.1156	8.6998	8.6478
South Korea (Won)	.0007734	.0007740	1293.00	1292.00
Spain (Peseta)	.005574	.005570	179.40	179.54
Sweden (Krona)	.0942	.0952	10.6178	10.5030
Switzerland (Franc)	.6303	.6260	1.5865	1.5974
1-month forward	.6302	.6259	1.5867	1.5977
3-months forward	.6301	.6258	1.5870	1.5980
6-months forward	.6303	.6259	1.5866	1.5977
Taiwan (Dollar)	.02894	.02895	34.560	34.540
Thailand (Baht)	.02266	.02265	44.140	44.150
Turkey (Lira)-f	.00000066	.00000066	1520000	1504000
United Arab (Dirham)	.2723	.2723	3.6730	3.6728
Uruguay (New Peso)				
Financial	.07278	.07278	13.740	13.740
Venezuela (Bolívar)	.001345	.001345	743.25	743.25
SDR	1.2945	1.2950	.7725	.7722
Euro	.9275	.9268	1.0782	1.0790

Special Drawing Rights (SDR) are based on exchange rates for the U.S., German, British, French, and Japanese currencies. Source: International Monetary Fund.
a-Russian Central Bank rate. b-Government rate. d-Floating rate; trading band suspended on 4/11/00. e-Adopted U.S. dollar as of 9/11/00. f-Floating rate, eff. Feb. 22.

Source: *Wall Street Journal*, Friday, September 21, 2001. Reprinted by permission of the *Wall Street Journal*, © 2001 Dow Jones & Company, Inc. All Rights Reserved Worldwide.

The Real Exchange Rate The **real exchange rate** is the relative price of the goods of two countries. That is, the real exchange rate tells us the rate at which we can trade the goods of one country for the goods of another. The real exchange rate is sometimes called the *terms of trade*.

To see the relation between the real and nominal exchange rates, consider a single good produced in many countries: cars. Suppose an American car costs \$10,000 and a similar Japanese car costs 2,400,000 yen. To compare the prices of the two cars, we must convert them into a common currency. If a dollar is worth 120 yen, then the American car costs 1,200,000 yen. Comparing the price of the American car (1,200,000 yen) and the price of the Japanese car (2,400,000 yen), we conclude that the American car costs one-half of what the Japanese car costs. In other words, at current prices, we can exchange two American cars for one Japanese car.

We can summarize our calculation as follows:

$$\begin{aligned}\text{Real Exchange Rate} &= \frac{(120 \text{ yen/dollar}) \times (10,000 \text{ dollars/American Car})}{(2,400,000 \text{ yen/Japanese Car})} \\ &= 0.5 \frac{\text{Japanese Car}}{\text{American Car}}.\end{aligned}$$

At these prices and this exchange rate, we obtain one-half of a Japanese car per American car. More generally, we can write this calculation as

$$\text{Real Exchange Rate} = \frac{\text{Nominal Exchange Rate} \times \text{Price of Domestic Good}}{\text{Price of Foreign Good}}.$$

The rate at which we exchange foreign and domestic goods depends on the prices of the goods in the local currencies and on the rate at which the currencies are exchanged.

This calculation of the real exchange rate for a single good suggests how we should define the real exchange rate for a broader basket of goods. Let e be the nominal exchange rate (the number of yen per dollar), P be the price level in the United States (measured in dollars), and P^* be the price level in Japan (measured in yen). Then the real exchange rate ϵ is

$$\begin{array}{ccccc}\text{Real} & & \text{Nominal} & & \text{Ratio of} \\ \text{Exchange} & = & \text{Exchange} & \times & \text{Price} \\ \text{Rate} & & \text{Rate} & & \text{Levels} \\ \epsilon & = & e & \times & (P/P^*).\end{array}$$

The real exchange rate between two countries is computed from the nominal exchange rate and the price levels in the two countries. *If the real exchange rate is high, foreign goods are relatively cheap, and domestic goods are relatively expensive. If the real exchange rate is low, foreign goods are relatively expensive, and domestic goods are relatively cheap.*

The Real Exchange Rate and the Trade Balance

What macroeconomic influence does the real exchange rate exert? To answer this question, remember that the real exchange rate is nothing more than a relative price. Just as the relative price of hamburgers and pizza determines which you choose for lunch, the relative price of domestic and foreign goods affects the demand for these goods.



Drawing by Matteson, © 1988
The New Yorker Magazine, Inc.

“How about Nebraska? The dollar’s still strong in Nebraska.”

Suppose first that the real exchange rate is low. In this case, because domestic goods are relatively cheap, domestic residents will want to purchase few imported goods: they will buy Fords rather than Toyotas, drink Coors rather than Heineken, and vacation in Florida rather than Europe. For the same reason, foreigners will want to buy many of our goods. As a result of both of these actions, the quantity of our net exports demanded will be high.


The opposite occurs if the real exchange rate is high. Because domestic goods are expensive relative to foreign goods, domestic residents will want to buy many imported goods, and foreigners will want to buy few of our goods. Therefore, the quantity of our net exports demanded will be low.

We write this relationship between the real exchange rate and net exports as

$$NX = NX(\epsilon).$$

This equation states that net exports are a function of the real exchange rate. Figure 5-7 illustrates this negative relationship between the trade balance and the real exchange rate.

figure 5-7



Net Exports and the Real Exchange Rate The figure shows the relationship between the real exchange rate and net exports: the lower the real exchange rate, the less expensive are domestic goods relative to foreign goods, and thus the greater are our net exports. Note that a portion of the horizontal axis measures negative values of NX : because imports can exceed exports, net exports can be less than zero.

The Determinants of the Real Exchange Rate


We now have all the pieces needed to construct a model that explains what factors determine the real exchange rate. In particular, we combine the relationship between net exports and the real exchange rate we just discussed with the model of the trade balance we developed earlier in the chapter. We can summarize the analysis as follows:

- The real exchange rate is related to net exports. When the real exchange rate is lower, domestic goods are less expensive relative to foreign goods, and net exports are greater.
- The trade balance (net exports) must equal the net capital outflow, which in turn equals saving minus investment. Saving is fixed by the consumption function and fiscal policy; investment is fixed by the investment function and the world interest rate.

Figure 5-8 illustrates these two conditions. The line showing the relationship between net exports and the real exchange rate slopes downward because a low real exchange rate makes domestic goods relatively inexpensive. The line representing the excess of saving over investment, $S - I$, is vertical because neither saving nor investment depends on the real exchange rate. The crossing of these two lines determines the equilibrium exchange rate.

Figure 5-8 looks like an ordinary supply-and-demand diagram. In fact, you can think of this diagram as representing the supply and demand for foreign-currency exchange. The vertical line, $S - I$, represents the net capital outflow and thus the supply of dollars to be exchanged into foreign currency and invested abroad. The downward-sloping line, NX , represents the net demand for dollars coming from foreigners who want dollars to buy our goods. *At the equilibrium real exchange rate, the supply of dollars available from the net capital outflow balances the demand for dollars by foreigners buying our net exports.*

figure 5-8



How the Real Exchange Rate Is Determined The real exchange rate is determined by the intersection of the vertical line representing saving minus investment and the downward-sloping net-exports schedule. At this intersection, the quantity of dollars supplied for the flow of capital abroad equals the quantity of dollars demanded for the net export of goods and services.


How Policies Influence the Real Exchange Rate

We can use this model to show how the changes in economic policy we discussed earlier affect the real exchange rate.

Fiscal Policy at Home What happens to the real exchange rate if the government reduces national saving by increasing government purchases or cutting taxes? As we discussed earlier, this reduction in saving lowers $S - I$ and thus NX . That is, the reduction in saving causes a trade deficit.

Figure 5-9 shows how the equilibrium real exchange rate adjusts to ensure that NX falls. The change in policy shifts the vertical $S - I$ line to the left, lowering the supply of dollars to be invested abroad. The lower supply causes the equilibrium real exchange rate to rise from ϵ_1 to ϵ_2 —that is, the dollar becomes more valuable. Because of the rise in the value of the dollar, domestic goods become more expensive relative to foreign goods, which causes exports to fall and imports to rise. The change in exports and the change in imports both act to reduce net exports.

figure 5-9




The Impact of Expansionary Fiscal Policy at Home on the Real Exchange Rate

Expansionary fiscal policy at home, such as an increase in government purchases or a cut in taxes, reduces national saving. The fall in saving reduces the supply of dollars to be exchanged into foreign currency, from $S_1 - I$ to $S_2 - I$. This shift raises the equilibrium real exchange rate from ϵ_1 to ϵ_2 .

Fiscal Policy Abroad What happens to the real exchange rate if foreign governments increase government purchases or cut taxes? This change in fiscal policy reduces world saving and raises the world interest rate. The increase in the world interest rate reduces domestic investment I , which raises $S - I$ and thus NX . That is, the increase in the world interest rate causes a trade surplus.

Figure 5-10 shows that this change in policy shifts the vertical $S - I$ line to the right, raising the supply of dollars to be invested abroad. The equilibrium real exchange rate falls. That is, the dollar becomes less valuable, and domestic goods become less expensive relative to foreign goods.

figure 5-10




The Impact of Expansionary Fiscal Policy Abroad on the Real Exchange Rate Expansionary fiscal policy abroad reduces world saving and raises the world interest rate from r_1^* to r_2^* . The increase in the world interest rate reduces investment at home, which in turn raises the supply of dollars to be exchanged into foreign currencies. As a result, the equilibrium real exchange rate falls from ϵ_1 to ϵ_2 .

Shifts in Investment Demand What happens to the real exchange rate if investment demand at home increases, perhaps because Congress passes an investment tax credit? At the given world interest rate, the increase in investment demand leads to higher investment. A higher value of I means lower values of $S - I$ and NX . That is, the increase in investment demand causes a trade deficit.

Figure 5-11 shows that the increase in investment demand shifts the vertical $S - I$ line to the left, reducing the supply of dollars to be invested abroad. The

figure 5-11



The Impact of an Increase in Investment Demand on the Real Exchange Rate An increase in investment demand raises the quantity of domestic investment from I_1 to I_2 . As a result, the supply of dollars to be exchanged into foreign currencies falls from $S - I_1$ to $S - I_2$. This fall in supply raises the equilibrium real exchange rate from ϵ_1 to ϵ_2 .


equilibrium real exchange rate rises. Hence, when the investment tax credit makes investing in the United States more attractive, it also increases the value of the U.S. dollars necessary to make these investments. When the dollar appreciates, domestic goods become more expensive relative to foreign goods, and net exports fall.

The Effects of Trade Policies

Now that we have a model that explains the trade balance and the real exchange rate, we have the tools to examine the macroeconomic effects of trade policies. Trade policies, broadly defined, are policies designed to influence directly the amount of goods and services exported or imported. Most often, trade policies take the form of protecting domestic industries from foreign competition—either by placing a tax on foreign imports (a tariff) or restricting the amount of goods and services that can be imported (a quota).

As an example of a protectionist trade policy, consider what would happen if the government prohibited the import of foreign cars. For any given real exchange rate, imports would now be lower, implying that net exports (exports minus imports) would be higher. Thus, the net-exports schedule shifts outward, as in Figure 5-12. To see the effects of the policy, we compare the old equilibrium and the new equilibrium. In the new equilibrium, the real exchange rate is higher, and net exports are unchanged. Despite the shift in the net-exports schedule, the equilibrium level of net exports remains the same, because the protectionist policy does not alter either saving or investment.

figure 5-12



The Impact of Protectionist Trade Policies on the Real Exchange Rate

A protectionist trade policy, such as a ban on imported cars, shifts the net-exports schedule from $NX(\epsilon)_1$ to $NX(\epsilon)_2$, which raises the real exchange rate from ϵ_1 to ϵ_2 . Notice that, despite the shift in the net-exports schedule, the equilibrium level of net exports is unchanged.

This analysis shows that protectionist trade policies do not affect the trade balance. This surprising conclusion is often overlooked in the popular debate over trade policies. Because a trade deficit reflects an excess of imports over exports, one might guess that reducing imports—such as by prohibiting the import of foreign cars—would reduce a trade deficit. Yet our model shows that protectionist policies lead only to an appreciation of the real exchange rate. The increase in the price of domestic goods relative to foreign goods tends to lower net exports by stimulating imports and depressing exports. Thus, the appreciation offsets the increase in net exports that is directly attributable to the trade restriction.

Although protectionist trade policies do not alter the trade balance, they do affect the amount of trade. As we have seen, because the real exchange rate appreciates, the goods and services we produce become more expensive relative to foreign goods and services. We therefore export less in the new equilibrium. Because net exports are unchanged, we must import less as well. (The appreciation of the exchange rate does stimulate imports to some extent, but this only partly offsets the decrease in imports caused by the trade restriction.) Thus, protectionist policies reduce both the quantity of imports and the quantity of exports.

This fall in the total amount of trade is the reason economists almost always oppose protectionist policies. International trade benefits all countries by allowing each country to specialize in what it produces best and by providing each country with a greater variety of goods and services. Protectionist policies diminish these gains from trade. Although these policies benefit certain groups within society—for example, a ban on imported cars helps domestic car producers—society on average is worse off when policies reduce the amount of international trade.

The Determinants of the Nominal Exchange Rate

Having seen what determines the real exchange rate, we now turn our attention to the nominal exchange rate—the rate at which the currencies of two countries trade. Recall the relationship between the real and the nominal exchange rate:

$$\begin{array}{ccccc} \text{Real} & & \text{Nominal} & & \text{Ratio of} \\ \text{Exchange} & = & \text{Exchange} & \times & \text{Price} \\ \text{Rate} & & \text{Rate} & & \text{Levels} \\ \epsilon & = & e & \times & (P/P^*). \end{array}$$

We can write the nominal exchange rate as

$$e = \epsilon \times (P^*/P).$$

This equation shows that the nominal exchange rate depends on the real exchange rate and the price levels in the two countries. Given the value of the real

exchange rate, if the domestic price level P rises, then the nominal exchange rate e will fall: because a dollar is worth less, a dollar will buy fewer yen. However, if the Japanese price level P^* rises, then the nominal exchange rate will increase: because the yen is worth less, a dollar will buy more yen.

It is instructive to consider changes in exchange rates over time. The exchange rate equation can be written

$$\% \text{ Change in } e = \% \text{ Change in } \epsilon + \% \text{ Change in } P^* - \% \text{ Change in } P.$$

The percentage change in ϵ is the change in the real exchange rate. The percentage change in P is the domestic inflation rate π , and the percentage change in P^* is the foreign country's inflation rate π^* . Thus, the percentage change in the nominal exchange rate is

$$\% \text{ Change in } e = \% \text{ Change in } \epsilon + (\pi^* - \pi)$$

$$\begin{array}{rcccl} \text{Percentage Change in} & = & \text{Percentage Change in} & + & \text{Difference in} \\ \text{Nominal Exchange Rate} & & \text{Real Exchange Rate} & & \text{Inflation Rates.} \end{array}$$

This equation states that the percentage change in the nominal exchange rate between the currencies of two countries equals the percentage change in the real exchange rate plus the difference in their inflation rates. *If a country has a high rate of inflation relative to the United States, a dollar will buy an increasing amount of the foreign currency over time. If a country has a low rate of inflation relative to the United States, a dollar will buy a decreasing amount of the foreign currency over time.*

This analysis shows how monetary policy affects the nominal exchange rate. We know from Chapter 4 that high growth in the money supply leads to high inflation. Here, we have just seen that one consequence of high inflation is a depreciating currency: high π implies falling e . In other words, just as growth in the amount of money raises the price of goods measured in terms of money, it also tends to raise the price of foreign currencies measured in terms of the domestic currency.

CASE STUDY

Inflation and Nominal Exchange Rates

If we look at data on exchange rates and price levels of different countries, we quickly see the importance of inflation for explaining changes in the nominal exchange rate. The most dramatic examples come from periods of very high inflation. For example, the price level in Mexico rose by 2,300 percent from 1983 to 1988. Because of this inflation, the number of pesos a person could buy with a U.S. dollar rose from 144 in 1983 to 2,281 in 1988.

The same relationship holds true for countries with more moderate inflation. Figure 5-13 is a scatterplot showing the relationship between inflation and the exchange rate for 15 countries. On the horizontal axis is the difference between each country's average inflation rate and the average inflation rate of the United

figure 5-13



Inflation Differentials and the Exchange Rate This scatterplot shows the relationship between inflation and the nominal exchange rate. The horizontal axis shows the country's average inflation rate minus the U.S. average inflation rate over the period 1972–2000. The vertical axis is the average percentage change in the country's exchange rate (per U.S. dollar) over that period. This figure shows that countries with relatively high inflation tend to have depreciating currencies and that countries with relatively low inflation tend to have appreciating currencies.

Source: International Financial Statistics.

States $(\pi^* - \pi)$. On the vertical axis is the average percentage change in the exchange rate between each country's currency and the U.S. dollar (% change in e). The positive relationship between these two variables is clear in this figure. Countries with relatively high inflation tend to have depreciating currencies (you can buy more of them for your dollars over time), and countries with relatively low inflation tend to have appreciating currencies (you can buy less of them for your dollars over time).

As an example, consider the exchange rate between German marks and U.S. dollars. Both Germany and the United States have experienced inflation over the past twenty years, so both the mark and the dollar buy fewer goods than they once did. But, as Figure 5-13 shows, inflation in Germany has been lower than inflation in the United States. This means that the value of the mark has fallen less than the value of the dollar. Therefore, the number of German marks you can buy with a U.S. dollar has been falling over time.


The Special Case of Purchasing-Power Parity

A famous hypothesis in economics, called the *law of one price*, states that the same good cannot sell for different prices in different locations at the same time. If a bushel of wheat sold for less in New York than in Chicago, it would be profitable to buy wheat in New York and then sell it in Chicago. Astute arbitrageurs would take advantage of such an opportunity and, thereby, would increase the demand for wheat in New York and increase the supply in Chicago. This would drive the price up in New York and down in Chicago—thereby ensuring that prices are equalized in the two markets.

The law of one price applied to the international marketplace is called **purchasing-power parity**. It states that if international arbitrage is possible, then a dollar (or any other currency) must have the same purchasing power in every country. The argument goes as follows. If a dollar could buy more wheat domestically than abroad, there would be opportunities to profit by buying wheat domestically and selling it abroad. Profit-seeking arbitrageurs would drive up the domestic price of wheat relative to the foreign price. Similarly, if a dollar could buy more wheat abroad than domestically, the arbitrageurs would buy wheat abroad and sell it domestically, driving down the domestic price relative to the foreign price. Thus, profit-seeking by international arbitrageurs causes wheat prices to be the same in all countries.

We can interpret the doctrine of purchasing-power parity using our model of the real exchange rate. The quick action of these international arbitrageurs implies that net exports are highly sensitive to small movements in the real exchange rate. A small decrease in the price of domestic goods relative to foreign goods—that is, a small decrease in the real exchange rate—causes arbitrageurs to buy goods domestically and sell them abroad. Similarly, a small increase in the relative price of domestic goods causes arbitrageurs to import goods from abroad. Therefore, as in Figure 5-14, the net-exports schedule is very flat at the real exchange rate that equalizes purchasing power among countries: any small

figure 5-14



Purchasing-Power Parity The law of one price applied to the international marketplace suggests that net exports are highly sensitive to small movements in the real exchange rate. This high sensitivity is reflected here with a very flat net-exports schedule.

movement in the real exchange rate leads to a large change in net exports. This extreme sensitivity of net exports guarantees that the equilibrium real exchange rate is always close to the level ensuring purchasing-power parity.

Purchasing-power parity has two important implications. First, because the net-exports schedule is flat, changes in saving or investment do not influence the real or nominal exchange rate. Second, because the real exchange rate is fixed, all changes in the nominal exchange rate result from changes in price levels.

Is this doctrine of purchasing-power parity realistic? Most economists believe that, despite its appealing logic, purchasing-power parity does not provide a completely accurate description of the world. First, many goods are not easily traded. A haircut can be more expensive in Tokyo than in New York, yet there is no room for international arbitrage because it is impossible to transport haircuts. Second, even tradable goods are not always perfect substitutes. Some consumers prefer Toyotas, and others prefer Fords. Thus, the relative price of Toyotas and Fords can vary to some extent without leaving any profit opportunities. For these reasons, real exchange rates do in fact vary over time.

Although the doctrine of purchasing-power parity does not describe the world perfectly, it does provide a reason why movement in the real exchange rate will be limited. There is much validity to its underlying logic: the farther the real exchange rate drifts from the level predicted by purchasing-power parity, the greater the incentive for individuals to engage in international arbitrage in goods. Although we cannot rely on purchasing-power parity to eliminate all changes in the real exchange rate, this doctrine does provide a reason to expect that fluctuations in the real exchange rate will typically be small or temporary.¹

CASE STUDY

The Big Mac Around the World

The doctrine of purchasing-power parity says that after we adjust for exchange rates, we should find that goods sell for the same price everywhere. Conversely, it says that the exchange rate between two currencies should depend on the price levels in the two countries.

To see how well this doctrine works, *The Economist*, an international news-magazine, regularly collects data on the price of a good sold in many countries: the McDonald's Big Mac hamburger. According to purchasing-power parity, the price of a Big Mac should be closely related to the country's nominal exchange rate. The higher the price of a Big Mac in the local currency, the higher the exchange rate (measured in units of local currency per U.S. dollar) should be.

Table 5-2 presents the international prices in 2000, when a Big Mac sold for \$2.51 in the United States. With these data we can use the doctrine of purchasing-power parity to predict nominal exchange rates. For example, because a Big Mac

¹ To learn more about purchasing-power parity, see Kenneth A. Froot and Kenneth Rogoff, "Perspectives on PPP and Long-Run Real Exchange Rates," in Gene M. Grossman and Kenneth Rogoff, eds., *Handbook of International Economics*, vol. 3 (Amsterdam: North-Holland, 1995).

table 5-2

**Big Mac Prices and the Exchange Rate:
An Application of Purchasing-Power Parity**

Country	Currency	Price of a Big Mac	Exchange Rate (per U.S. dollar)	
			Predicted	Actual
Indonesia	Rupiah	14,500	5,777	7,945
Italy	Lira	4,500	1,793	2,088
South Korea	Won	3,000	1,195	1,108
Chile	Peso	1,260	502	514
Spain	Peseta	375	149	179
Hungary	Forint	339	135	279
Japan	Yen	294	117	106
Taiwan	Dollar	70.0	27.9	30.6
Thailand	Baht	55.0	21.9	38.0
Czech Rep.	Crown	54.37	21.7	39.1
Russia	Ruble	39.50	15.7	28.5
Denmark	Crown	24.75	9.86	8.04
Sweden	Crown	24.0	9.56	8.84
Mexico	Peso	20.9	8.33	9.41
France	Franc	18.5	7.37	7.07
Israel	Shekel	14.5	5.78	4.05
China	Yuan	9.90	3.94	8.28
South Africa	Rand	9.0	3.59	6.72
Switzerland	Franc	5.90	2.35	1.70
Poland	Zloty	5.50	2.19	4.30
Germany	Mark	4.99	1.99	2.11
Malaysia	Dollar	4.52	1.80	3.80
New Zealand	Dollar	3.40	1.35	2.01
Singapore	Dollar	3.20	1.27	1.70
Brazil	Real	2.95	1.18	1.79
Canada	Dollar	2.85	1.14	1.47
Australia	Dollar	2.59	1.03	1.68
United States	Dollar	2.51	1.00	1.00
Argentina	Peso	2.50	1.00	1.00
Britain	Pound	1.90	0.76	0.63

Note: The predicted exchange rate is the exchange rate that would make the price of a Big Mac in that country equal to its price in the United States.

Source: *The Economist*, April 29, 2000, 75.

cost 294 yen in Japan, we would predict that the exchange rate between the dollar and the yen was $294/2.51$, or 117, yen per dollar. At this exchange rate, a Big Mac would have cost the same in Japan and the United States.

Table 5-2 shows the predicted and actual exchange rates for 30 countries, ranked by the predicted exchange rate. You can see that the evidence on purchasing-power

parity is mixed. As the last two columns show, the actual and predicted exchange rate are usually in the same ballpark. Our theory predicts, for instance, that a U.S. dollar should buy the greatest number of Indonesian rupiahs and fewest British pounds, and this turns out to be true. In the case of Japan, the predicted exchange rate of 117 yen per dollar is close to the actual exchange rate of 106. Yet the theory's predictions are far from exact and, in many cases, are off by 30 percent or more. Hence, although the theory of purchasing-power parity provides a rough guide to the level of exchange rates, it does not explain exchange rates completely.

5-4 Conclusion: The United States as a Large Open Economy

In this chapter we have seen how a small open economy works. We have examined the determinants of the international flow of funds for capital accumulation and the international flow of goods and services. We have also examined the determinants of a country's real and nominal exchange rates. Our analysis shows how various policies—monetary policies, fiscal policies, and trade policies—affect the trade balance and the exchange rate.

The economy we have studied is “small” in the sense that its interest rate is fixed by world financial markets. That is, we have assumed that this economy does not affect the world interest rate, and that the economy can borrow and lend at the world interest rate in unlimited amounts. This assumption contrasts with the assumption we made when we studied the closed economy in Chapter 3. In the closed economy, the domestic interest rate equilibrates domestic saving and domestic investment, implying that policies that influence saving or investment alter the equilibrium interest rate.

Which of these analyses should we apply to an economy such as the United States? The answer is a little of both. The United States is neither so large nor so isolated that it is immune to developments occurring abroad. The large trade deficits of the 1980s and 1990s show the importance of international financial markets for funding U.S. investment. Hence, the closed-economy analysis of Chapter 3 cannot by itself fully explain the impact of policies on the U.S. economy.

Yet the U.S. economy is not so small and so open that the analysis of this chapter applies perfectly either. First, the United States is large enough that it can influence world financial markets. For example, large U.S. budget deficits were often blamed for the high real interest rates that prevailed throughout the world in the 1980s. Second, capital may not be perfectly mobile across countries. If individuals prefer holding their wealth in domestic rather than foreign assets, funds for capital accumulation will not flow freely to equate interest rates in all countries. For these two reasons, we cannot directly apply our model of the small open economy to the United States.

When analyzing policy for a country such as the United States, we need to combine the closed-economy logic of Chapter 3 and the small-open-economy logic of this chapter. The appendix to this chapter builds a model of an economy between these two extremes. In this intermediate case, there is international borrowing and lending, but the interest rate is not fixed by world financial markets. Instead, the more the economy borrows from abroad, the higher the interest rate it must offer foreign investors. The results, not surprisingly, are a mixture of the two polar cases we have already examined.

Consider, for example, a reduction in national saving caused by a fiscal expansion. As in the closed economy, this policy raises the real interest rate and crowds out domestic investment. As in the small open economy, it also reduces the net capital outflow, leading to a trade deficit and an appreciation of the exchange rate. Hence, although the model of the small open economy examined here does not precisely describe an economy such as the United States, it does provide approximately the right answer to how policies affect the trade balance and the exchange rate.

Summary

1. Net exports are the difference between exports and imports. They are equal to the difference between what we produce and what we demand for consumption, investment, and government purchases.
2. The net capital outflow is the excess of domestic saving over domestic investment. The trade balance is the amount received for our net exports of goods and services. The national income accounts identity shows that the net capital outflow always equals the trade balance.
3. The impact of any policy on the trade balance can be determined by examining its impact on saving and investment. Policies that raise saving or lower investment lead to a trade surplus, and policies that lower saving or raise investment lead to a trade deficit.
4. The nominal exchange rate is the rate at which people trade the currency of one country for the currency of another country. The real exchange rate is the rate at which people trade the goods produced by the two countries. The real exchange rate equals the nominal exchange rate multiplied by the ratio of the price levels in the two countries.
5. Because the real exchange rate is the price of domestic goods relative to foreign goods, an appreciation of the real exchange rate tends to reduce net exports. The equilibrium real exchange rate is the rate at which the quantity of net exports demanded equals the net capital outflow.
6. The nominal exchange rate is determined by the real exchange rate and the price levels in the two countries. Other things equal, a high rate of inflation leads to a depreciating currency.

KEY CONCEPTS

Net exports	Balanced trade	Nominal exchange rate
Trade balance	Small open economy	Real exchange rate
Net capital outflow	World interest rate	Purchasing-power parity
Trade surplus and trade deficit		

QUESTIONS FOR REVIEW

1. What are the net capital outflow and the trade balance? Explain how they are related.
2. Define the nominal exchange rate and the real exchange rate.
3. If a small open economy cuts defense spending, what happens to saving, investment, the trade balance, the interest rate, and the exchange rate?
4. If a small open economy bans the import of Japanese VCRs, what happens to saving, investment, the trade balance, the interest rate, and the exchange rate?
5. If Germany has low inflation and Italy has high inflation, what will happen to the exchange rate between the German mark and the Italian lira?

PROBLEMS AND APPLICATIONS

1. Use the model of the small open economy to predict what would happen to the trade balance, the real exchange rate, and the nominal exchange rate in response to each of the following events.
 - a. A fall in consumer confidence about the future induces consumers to spend less and save more.
 - b. The introduction of a stylish line of Toyotas makes some consumers prefer foreign cars over domestic cars.
 - c. The introduction of automatic teller machines reduces the demand for money.
2. Consider an economy described by the following equations:

$$Y = C + I + G + NX,$$

$$Y = 5,000,$$

$$G = 1,000,$$

$$T = 1,000,$$

$$C = 250 + 0.75(Y - T),$$

$$I = 1,000 - 50r,$$

$$NX = 500 - 500\epsilon,$$

$$r = r^* = 5.$$
 - a. In this economy, solve for national saving, investment, the trade balance, and the equilibrium exchange rate.
 - b. Suppose now that G rises to 1,250. Solve for national saving, investment, the trade balance, and the equilibrium exchange rate. Explain what you find.
 - c. Now suppose that the world interest rate rises from 5 to 10 percent. (G is again 1,000). Solve for national saving, investment, the trade balance, and the equilibrium exchange rate. Explain what you find.
3. The country of Leverett is a small open economy. Suddenly, a change in world fashions makes the exports of Leverett unpopular.
 - a. What happens in Leverett to saving, investment, net exports, the interest rate, and the exchange rate?
 - b. The citizens of Leverett like to travel abroad. How will this change in the exchange rate affect them?
 - c. The fiscal policymakers of Leverett want to adjust taxes to maintain the exchange rate at

its previous level. What should they do? If they do this, what are the overall effects on saving, investment, net exports, and the interest rate?

4. What will happen to the trade balance and the real exchange rate of a small open economy when government purchases increase, such as during a war? Does your answer depend on whether this is a local war or a world war?
5. In 1995, President Clinton considered placing a 100-percent tariff on the import of Japanese luxury cars. Discuss the economics and politics of such a policy. In particular, how would the policy affect the U.S. trade deficit? How would it affect the exchange rate? Who would be hurt by such a policy? Who would benefit?
6. Suppose that some foreign countries begin to subsidize investment by instituting an investment tax credit.
 - a. What happens to world investment demand as a function of the world interest rate?
 - b. What happens to the world interest rate?
 - c. What happens to investment in our small open economy?
 - d. What happens to our trade balance?
 - e. What happens to our real exchange rate?
7. “Traveling in Italy is much cheaper now than it was ten years ago,” says a friend. “Ten years ago, a dollar bought 1,000 lire; this year, a dollar buys 1,500 lire.”

Is your friend right or wrong? Given that total inflation over this period was 25 percent in the United States and 100 percent in Italy, has it become more or less expensive to travel in Italy? Write your answer using a concrete example—such as a cup of American coffee versus a cup of Italian espresso—that will convince your friend.
8. You read in a newspaper that the nominal interest rate is 12 percent per year in Canada and 8 percent per year in the United States. Suppose that the real interest rates are equalized in the two countries and that purchasing-power parity holds.
 - a. Using the Fisher equation (discussed in Chapter 4), what can you infer about expected inflation in Canada and in the United States?
 - b. What can you infer about the expected change in the exchange rate between the Canadian dollar and the U.S. dollar?
 - c. A friend proposes a get-rich-quick scheme: borrow from a U.S. bank at 8 percent, deposit the money in a Canadian bank at 12 percent, and make a 4 percent profit. What’s wrong with this scheme?

APPENDIX

The Large Open Economy

When analyzing policy for a country such as the United States, we need to combine the closed-economy logic of Chapter 3 and the small-open-economy logic of this chapter. This appendix presents a model of an economy between these two extremes, called the *large open economy*.

Net Capital Outflow

The key difference between the small and large open economies is the behavior of the net capital outflow. In the model of the small open economy, capital flows freely into or out of the economy at a fixed world interest rate r^* . The model of the large open economy makes a different assumption about international capital flows. To understand that assumption, keep in mind that the net capital outflow is the amount that domestic investors lend abroad minus the amount that foreign investors lend here.

Imagine that you are a domestic investor—such as the portfolio manager of a university endowment—deciding where to invest your funds. You could invest domestically (for example, by making loans to U.S. companies), or you could invest abroad (by making loans to foreign companies). Many factors may affect your decision, but surely one of them is the interest rate you can earn. The higher the interest rate you can earn domestically, the less attractive you would find foreign investment.

Investors abroad face a similar decision. They have a choice between investing in their home country or lending to someone in the United States. The higher the interest rate in the United States, the more willing foreigners are to lend to U.S. companies and to buy U.S. assets.


Thus, because of the behavior of both domestic and foreign investors, the net flow of capital to other countries, which we'll denote as CF , is negatively related to the domestic real interest rate r . As the interest rate rises, less of our saving flows abroad, and more funds for capital accumulation flow in from other countries. We write this as

$$CF = CF(r).$$

This equation states that the net capital outflow is a function of the domestic interest rate. Figure 5-15 on page 146 illustrates this relationship. Notice that CF can be either positive or negative, depending on whether the economy is a lender or borrower in world financial markets.

To see how this CF function relates to our previous models, consider Figure 5-16 on page 146. This figure shows two special cases: a vertical CF function and a horizontal CF function.

figure 5-15




How the Net Capital Outflow Depends on the Interest Rate A higher domestic interest rate discourages domestic investors from lending abroad and encourages foreign investors to lend here. Therefore, net capital outflow CF is negatively related to the interest rate.

The closed economy is the special case shown in panel (a) of Figure 5-16. In the closed economy, there is no international borrowing or lending, and the interest rate adjusts to equilibrate domestic saving and investment. This means that $CF = 0$ at all interest rates. This situation would arise if investors here and abroad were unwilling to hold foreign assets, regardless of the return. It might also arise if the government prohibited its citizens from transacting in foreign financial markets, as some governments do.

The small open economy with perfect capital mobility is the special case shown in panel (b) of Figure 5-16. In this case, capital flows freely into and out of the country at the fixed world interest rate r^* . This situation would arise if investors here and abroad bought whatever asset yielded the highest

figure 5-16



Two Special Cases In the closed economy, shown in panel (a), the net capital outflow is zero for all interest rates. In the small open economy with perfect capital mobility, shown in panel (b), the net capital outflow is perfectly elastic at the world interest rate r^* .

return, and if this economy were too small to affect the world interest rate. The economy's interest rate would be fixed at the interest rate prevailing in world financial markets.

Why isn't the interest rate of a large open economy such as the United States fixed by the world interest rate? There are two reasons. The first is that the United States is large enough to influence world financial markets. The more the United States lends abroad, the greater the supply of loans in the world economy is, and the lower interest rates become around the world. The more the United States borrows from abroad (that is, the more negative CF becomes), the higher world interest rates are. We use the label "large open economy" because this model applies to an economy large enough to affect world interest rates.

There is, however, a second reason that the interest rate in an economy may not be fixed by the world interest rate: capital may not be perfectly mobile. That is, investors here and abroad may prefer to hold their wealth in domestic rather than foreign assets. Such a preference for domestic assets could arise because of imperfect information about foreign assets or because of government impediments to international borrowing and lending. In either case, funds for capital accumulation will not flow freely to equalize interest rates in all countries. Instead, the net capital outflow will depend on domestic interest rates relative to foreign interest rates. U.S. investors will lend abroad only if U.S. interest rates are comparatively low, and foreign investors will lend in the United States only if U.S. interest rates are comparatively high. The large-open-economy model, therefore, may apply even to a small economy if capital does not flow freely into and out of the economy.

Hence, either because the large open economy affects world interest rates, or because capital is imperfectly mobile, or perhaps for both reasons, the CF function slopes downward. Except for this new downward-sloping CF function, the model of the large open economy resembles the model of the small open economy. We put all the pieces together in the next section.

The Model


To understand how the large open economy works, we need to consider two key markets: the market for loanable funds (where the interest rate is determined) and the market for foreign exchange (where the exchange rate is determined). The interest rate and the exchange rate are two prices that guide the allocation of resources.

The Market for Loanable Funds An open economy's saving S is used in two ways: to finance domestic investment I and to finance the net capital outflow CF . We can write

$$S = I + CF.$$

Consider how these three variables are determined. National saving is fixed by the level of output, fiscal policy, and the consumption function. Investment

figure 5-17



The Market for Loanable Funds in the Large Open Economy At the equilibrium interest rate, the supply of loanable funds from saving S balances the demand for loanable funds from domestic investment I and capital investments abroad CF .

and net capital outflow both depend on the domestic real interest rate. We can write

$$\bar{S} = I(r) + CF(r).$$

Figure 5-17 shows the market for loanable funds. The supply of loanable funds is national saving. The demand for loanable funds is the sum of the demand for domestic investment and the demand for foreign investment (net capital outflow). The interest rate adjusts to equilibrate supply and demand.

The Market for Foreign Exchange Next, consider the relationship between the net capital outflow and the trade balance. The national income accounts identity tells us


$$NX = S - I.$$

Because NX is a function of the real exchange rate, and because $CF = S - I$, we can write

$$NX(\epsilon) = CF.$$

Figure 5-18 shows the equilibrium in the market for foreign exchange. Once again, the real exchange rate is the price that equilibrates the trade balance and the net capital outflow.

figure 5-18



The Market for Foreign-Currency Exchange in the Large Open Economy At the equilibrium exchange rate, the supply of dollars from the net capital outflow, CF , balances the demand for dollars from our net exports of goods and services, NX .

The last variable we should consider is the nominal exchange rate. As before, the nominal exchange rate is the real exchange rate times the ratio of the price levels:

$$e = \epsilon \times (P^*/P).$$

The real exchange rate is determined as in Figure 5-18, and the price levels are determined by monetary policies here and abroad, as we discussed in Chapter 4. Forces that move the real exchange rate or the price levels also move the nominal exchange rate.

Policies in the Large Open Economy

We can now consider how economic policies influence the large open economy. Figure 5-19 shows the three diagrams we need for the analysis. Panel (a) shows the equilibrium in the market for loanable funds; panel (b) shows the relationship between the equilibrium interest rate and the net capital outflow; and panel (c) shows the equilibrium in the market for foreign exchange.

Fiscal Policy at Home Consider the effects of expansionary fiscal policy—an increase in government purchases or a decrease in taxes. Figure 5-20 shows what

figure 5-19





figure 5-20



A Reduction in National Saving in the Large Open Economy Panel (a) shows that a reduction in national saving lowers the supply of loanable funds. The equilibrium interest rate rises. Panel (b) shows that the higher interest rate lowers the net capital outflow. Panel (c) shows that the reduced capital outflow means a reduced supply of dollars in the market for foreign-currency exchange. The reduced supply of dollars causes the real exchange rate to appreciate and net exports to fall.

happens. The policy reduces national saving S , thereby reducing the supply of loanable funds and raising the equilibrium interest rate r . The higher interest rate reduces both domestic investment I and the net capital outflow CF . The fall in the net capital outflow reduces the supply of dollars to be exchanged into foreign currency. The exchange rate appreciates, and net exports fall.

Note that the impact of fiscal policy in this model combines its impact in the closed economy and its impact in the small open economy. As in the closed economy, a fiscal expansion in a large open economy raises the interest rate and crowds out investment. As in the small open economy, a fiscal expansion causes a trade deficit and an appreciation in the exchange rate.

One way to see how the three types of economy are related is to consider the identity


$$S = I + NX.$$

In all three cases, expansionary fiscal policy reduces national saving S . In the closed economy, the fall in S coincides with an equal fall in I , and NX stays constant at zero. In the small open economy, the fall in S coincides with an equal fall

in NX , and I remains constant at the level fixed by the world interest rate. The large open economy is the intermediate case: both I and NX fall, each by less than the fall in S .


Shifts in Investment Demand Suppose that the investment demand schedule shifts outward, perhaps because Congress passes an investment tax credit. Figure 5-21 shows the effect. The demand for loanable funds rises, raising the equilibrium interest rate. The higher interest rate reduces the net capital outflow: Americans make fewer loans abroad, and foreigners make more loans to Americans. The fall in the net capital outflow reduces the supply of dollars in the market for foreign exchange. The exchange rate appreciates, and net exports fall.

figure 5-21



Trade Policies Figure 5-22 shows the effect of a trade restriction, such as an import quota. The reduced demand for imports shifts the net-exports schedule outward in panel (c). Because nothing has changed in the market for loanable funds, the interest rate remains the same, which in turn implies that the net capital outflow remains the same. The shift in the net-exports schedule causes the exchange rate to appreciate. The rise in the exchange rate makes U.S. goods

figure 5-22




expensive relative to foreign goods, which depresses exports and stimulates imports. In the end, the trade restriction does not affect the trade balance.

Shifts in Net Capital Outflow There are various reasons that the CF schedule might shift. One reason is fiscal policy abroad. For example, suppose that Germany pursues a fiscal policy that raises German saving. This policy reduces the German interest rate. The lower German interest rate discourages American investors from lending in Germany and encourages German investors to lend in the United States. For any given U.S. interest rate, the U.S. net capital outflow falls.

Another reason the CF schedule might shift is political instability abroad. Suppose that a war or revolution breaks out in another country. Investors around the world will try to withdraw their assets from that country and seek a “safe haven” in a stable country such as the United States. The result is a reduction in the U.S. net capital outflow.

Figure 5-23 shows the impact of a shift in the CF schedule. The reduced demand for loanable funds lowers the equilibrium interest rate. The lower interest

figure 5-23



rate tends to raise net capital outflow, but because this only partly mitigates the shift in the CF schedule, CF still falls. The reduced level of net capital outflow reduces the supply of dollars in the market for foreign exchange. The exchange rate appreciates, and net exports fall.

Conclusion

How different are large and small open economies? Certainly, policies affect the interest rate in a large open economy, unlike in a small open economy. But, in other ways, the two models yield similar conclusions. In both large and small open economies, policies that raise saving or lower investment lead to trade surpluses. Similarly, policies that lower saving or raise investment lead to trade deficits. In both economies, protectionist trade policies cause the exchange rate to appreciate and do not influence the trade balance. Because the results are so similar, for most questions one can use the simpler model of the small open economy, even if the economy being examined is not really small.

MORE PROBLEMS AND APPLICATIONS

1. If a war broke out abroad, it would affect the U.S. economy in many ways. Use the model of the large open economy to examine each of the following effects of such a war. What happens in the United States to saving, investment, the trade balance, the interest rate, and the exchange rate? (To keep things simple, consider each of the following effects separately.)
 - a. The U.S. government, fearing it may need to enter the war, increases its purchases of military equipment.
 - b. Other countries raise their demand for high-tech weapons, a major export of the United States.
 - c. The war makes U.S. firms uncertain about the future, and the firms delay some investment projects.
 - d. The war makes U.S. consumers uncertain about the future, and the consumers save more in response.
 - e. Americans become apprehensive about traveling abroad, so more of them spend their vacations in the United States.
 - f. Foreign investors seek a safe haven for their portfolios in the United States.
2. On September 21, 1995, “House Speaker Newt Gingrich threatened to send the United States into default on its debt for the first time in the nation’s history, to force the Clinton Administration to balance the budget on Republican terms” (*New York Times*, September 22, 1995, A1). That same day, the interest rate on 30-year U.S. government bonds rose from 6.46 to 6.55 percent, and the dollar fell in value from 102.7 to 99.0 yen. Use the model of the large open economy to explain this event.

Unemployment

A man willing to work, and unable to find work, is perhaps the saddest sight that fortune's inequality exhibits under the sun.

—Thomas Carlyle

Unemployment is the macroeconomic problem that affects people most directly and severely. For most people, the loss of a job means a reduced living standard and psychological distress. It is no surprise that unemployment is a frequent topic of political debate and that politicians often claim that their proposed policies would help create jobs.


Economists study unemployment to identify its causes and to help improve the public policies that affect the unemployed. Some of these policies, such as job-training programs, assist people in finding employment. Others, such as unemployment insurance, alleviate some of the hardships that the unemployed face. Still other policies affect the prevalence of unemployment inadvertently. Laws mandating a high minimum wage, for instance, are widely thought to raise unemployment among the least skilled and experienced members of the labor force. By showing the effects of various policies, economists help policymakers evaluate their options.

Our discussions of the labor market so far have ignored unemployment. In particular, the model of national income in Chapter 3 was built with the assumption that the economy was always at full employment. In reality, of course, not everyone in the labor force has a job all the time: all free-market economies experience some unemployment.

Figure 6-1 shows the rate of unemployment—the percentage of the labor force unemployed—in the United States since 1948. Although the rate of unemployment fluctuates from year to year, it never gets even close to zero. The average is between 5 and 6 percent, meaning that about 1 out of every 18 people wanting a job does not have one.

In this chapter we begin our study of unemployment by discussing why there is always some unemployment and what determines its level. We do not study what determines the year-to-year fluctuations in the rate of unemployment until Part IV of this book, where we examine short-run economic fluctuations. Here we examine the determinants of the **natural rate of unemployment**—the average rate of unemployment around which the economy fluctuates. The natural

figure 6-1



The Unemployment Rate and the Natural Rate of Unemployment in the United States There is always some unemployment. The natural rate of unemployment is the average level around which the unemployment rate fluctuates. (The natural rate of unemployment for any particular year is estimated here by averaging all the unemployment rates from ten years earlier to ten years later. Future unemployment rates are set at 5.5 percent.)

rate is the rate of unemployment toward which the economy gravitates in the long run, given all the labor-market imperfections that impede workers from instantly finding jobs.

6-1 Job Loss, Job Finding, and the Natural Rate of Unemployment

Every day some workers lose or quit their jobs, and some unemployed workers are hired. This perpetual ebb and flow determines the fraction of the labor force that is unemployed. In this section we develop a model of labor-force dynamics that shows what determines the natural rate of unemployment.¹

We start with some notation. Let L denote the labor force, E the number of employed workers, and U the number of unemployed workers. Because every

¹ Robert E. Hall, "A Theory of the Natural Rate of Unemployment and the Duration of Unemployment," *Journal of Monetary Economics* 5 (April 1979): 153–169.

worker is either employed or unemployed, the labor force is the sum of the employed and the unemployed:

$$L = E + U.$$

In this notation, the rate of unemployment is U/L .

To see what determines the unemployment rate, we assume that the labor force L is fixed and focus on the transition of individuals in the labor force between employment and unemployment. This is illustrated in Figure 6-2. Let s denote the rate of job separation, the fraction of employed individuals who lose their job each month. Let f denote the rate of job finding, the fraction of unemployed individuals who find a job each month. Together, the rate of job separation s and the rate of job finding f determine the rate of unemployment.


If the unemployment rate is neither rising nor falling—that is, if the labor market is in a steady state—then the number of people finding jobs must equal the number of people losing jobs. The number of people finding jobs is fU and the number of people losing jobs is sE , so we can write the steady-state condition as

$$fU = sE.$$

We can use this equation to find the steady-state unemployment rate. From an earlier equation, we know that $E = L - U$; that is, the number of employed equals the labor force minus the number of unemployed. If we substitute $(L - U)$ for E in the steady-state condition, we find

$$fU = s(L - U).$$

figure 6-2



The Transitions Between Employment and Unemployment In every period, a fraction s of the employed lose their jobs, and a fraction f of the unemployed find jobs. The rates of job separation and job finding determine the rate of unemployment.

To get closer to solving for the unemployment rate, divide both sides of this equation by L to obtain

$$f \frac{U}{L} = s(1 - \frac{U}{L}).$$

Now we can solve for U/L to find

$$\frac{U}{L} = \frac{s}{s + f}.$$

This equation shows that the steady-state rate of unemployment U/L depends on the rates of job separation s and job finding f . The higher the rate of job separation, the higher the unemployment rate. The higher the rate of job finding, the lower the unemployment rate.

Here's a numerical example. Suppose that 1 percent of the employed lose their jobs each month ($s = 0.01$). This means that on average jobs last 100 months, or about 8 years. Suppose further that about 20 percent of the unemployed find a job each month ($f = 0.20$), so that spells of unemployment last 5 months on average. Then the steady-state rate of unemployment is

$$\begin{aligned} \frac{U}{L} &= \frac{0.01}{0.01 + 0.20} \\ &= 0.0476. \end{aligned}$$

The rate of unemployment in this example is about 5 percent.

This model of the natural rate of unemployment has an obvious but important implication for public policy. *Any policy aimed at lowering the natural rate of unemployment must either reduce the rate of job separation or increase the rate of job finding. Similarly, any policy that affects the rate of job separation or job finding also changes the natural rate of unemployment.*

Although this model is useful in relating the unemployment rate to job separation and job finding, it fails to answer a central question: Why is there unemployment in the first place? If a person could always find a job quickly, then the rate of job finding would be very high and the rate of unemployment would be near zero. This model of the unemployment rate assumes that job finding is not instantaneous, but it fails to explain why. In the next two sections, we examine two underlying reasons for unemployment: job search and wage rigidity.

6-2 Job Search and Frictional Unemployment

One reason for unemployment is that it takes time to match workers and jobs. The equilibrium model of the aggregate labor market discussed in Chapter 3 assumes that all workers and all jobs are identical, and therefore that all workers are

equally well suited for all jobs. If this were true and the labor market were in equilibrium, then a job loss would not cause unemployment: a laid-off worker would immediately find a new job at the market wage.

In fact, workers have different preferences and abilities, and jobs have different attributes. Furthermore, the flow of information about job candidates and job vacancies is imperfect, and the geographic mobility of workers is not instantaneous. For all these reasons, searching for an appropriate job takes time and effort, and this tends to reduce the rate of job finding. Indeed, because different jobs require different skills and pay different wages, unemployed workers may not accept the first job offer they receive. The unemployment caused by the time it takes workers to search for a job is called **frictional unemployment**.

Some frictional unemployment is inevitable in a changing economy. For many reasons, the types of goods that firms and households demand vary over time. As the demand for goods shifts, so does the demand for the labor that produces those goods. The invention of the personal computer, for example, reduced the demand for typewriters and, as a result, for labor by typewriter manufacturers. At the same time, it increased the demand for labor in the electronics industry. Similarly, because different regions produce different goods, the demand for labor may be rising in one part of the country and falling in another. A decline in the price of oil may cause the demand for labor to fall in oil-producing states such as Texas, but because cheap oil makes driving less expensive, it increases the demand for labor in auto-producing states such as Michigan. Economists call a change in the composition of demand among industries or regions a **sectoral shift**. Because sectoral shifts are always occurring, and because it takes time for workers to change sectors, there is always frictional unemployment.

Sectoral shifts are not the only cause of job separation and frictional unemployment. In addition, workers find themselves unexpectedly out of work when their firms fail, when their job performance is deemed unacceptable, or when their particular skills are no longer needed. Workers also may quit their jobs to change careers or to move to different parts of the country. As long as the supply and demand for labor among firms is changing, frictional unemployment is unavoidable.

Public Policy and Frictional Unemployment

Many public policies seek to decrease the natural rate of unemployment by reducing frictional unemployment. Government employment agencies disseminate information about job vacancies in order to match jobs and workers more efficiently. Publicly funded retraining programs are designed to ease the transition of workers from declining to growing industries. If these programs succeed at increasing the rate of job finding, they decrease the natural rate of unemployment.

Other government programs inadvertently increase the amount of frictional unemployment. One of these is **unemployment insurance**. Under

this program, unemployed workers can collect a fraction of their wages for a certain period after losing their jobs. Although the precise terms of the program differ from year to year and from state to state, a typical worker covered by unemployment insurance in the United States receives 50 percent of his or her former wages for 26 weeks. In many European countries, unemployment-insurance programs are even more generous.

By softening the economic hardship of unemployment, unemployment insurance increases the amount of frictional unemployment and raises the natural rate. The unemployed who receive unemployment-insurance benefits are less pressed to search for new employment and are more likely to turn down unattractive job offers. Both of these changes in behavior reduce the rate of job finding. In addition, because workers know that their incomes are partially protected by unemployment insurance, they are less likely to seek jobs with stable employment prospects and are less likely to bargain for guarantees of job security. These behavioral changes raise the rate of job separation.

That unemployment insurance raises the natural rate of unemployment does not necessarily imply that the policy is ill advised. The program has the benefit of reducing workers' uncertainty about their incomes. Moreover, inducing workers to reject unattractive job offers may lead to a better matching between workers and jobs. Evaluating the costs and benefits of different systems of unemployment insurance is a difficult task that continues to be a topic of much research.

Economists who study unemployment insurance often propose reforms that would reduce the amount of unemployment. One common proposal is to require a firm that lays off a worker to bear the full cost of that worker's unemployment benefits. Such a system is called *100 percent experience rated*, because the rate that each firm pays into the unemployment-insurance system fully reflects the unemployment experience of its own workers. Most current programs are *partially experience rated*. Under this system, when a firm lays off a worker, it is charged for only part of the worker's unemployment benefits; the remainder comes from the program's general revenue. Because a firm pays only a fraction of the cost of the unemployment it causes, it has an incentive to lay off workers when its demand for labor is temporarily low. By reducing that incentive, the proposed reform may reduce the prevalence of temporary layoffs.

CASE STUDY

Unemployment Insurance and the Rate of Job Finding

Many studies have examined the effect of unemployment insurance on job search. The most persuasive studies use data on the experiences of unemployed individuals, rather than economy-wide rates of unemployment. Individual data often yield sharp results that are open to few alternative explanations.

One study followed the experience of individual workers as they used up their eligibility for unemployment-insurance benefits. It found that when

unemployed workers become ineligible for benefits, they are more likely to find new jobs. In particular, the probability of a person finding a new job more than doubles when his or her benefits run out. One possible explanation is that an absence of benefits increases the search effort of unemployed workers. Another possibility is that workers without benefits are more likely to accept job offers that would otherwise be declined because of low wages or poor working conditions.²

Additional evidence on how economic incentives affect job search comes from an experiment that the state of Illinois ran in 1985. Randomly selected new claimants for unemployment insurance were each offered a \$500 bonus if they found employment within 11 weeks. The subsequent experience of this group was compared to that of a control group not offered the incentive. The average duration of unemployment for the group offered the \$500 bonus was 17.0 weeks, compared to 18.3 weeks for the control group. Thus, the bonus reduced the average spell of unemployment by 7 percent, suggesting that more effort was devoted to job search. This experiment shows clearly that the incentives provided by the unemployment-insurance system affect the rate of job finding.³

6-3 Real-Wage Rigidity and Structural Unemployment

A second reason for unemployment is **wage rigidity**—the failure of wages to adjust until labor supply equals labor demand. In the equilibrium model of the labor market, as outlined in Chapter 3, the real wage adjusts to equilibrate supply and demand. Yet wages are not always flexible. Sometimes the real wage is stuck above the market-clearing level.


Figure 6-3 shows why wage rigidity leads to unemployment. When the real wage is above the level that equilibrates supply and demand, the quantity of labor supplied exceeds the quantity demanded. Firms must in some way ration the scarce jobs among workers. Real-wage rigidity reduces the rate of job finding and raises the level of unemployment.

The unemployment resulting from wage rigidity and job rationing is called **structural unemployment**. Workers are unemployed not because they are actively searching for the jobs that best suit their individual skills but because, at the going wage, the supply of labor exceeds the demand. These workers are simply waiting for jobs to become available.

² Lawrence F. Katz and Bruce D. Meyer, “Unemployment Insurance, Recall Expectations, and Unemployment Outcomes,” *Quarterly Journal of Economics* 105 (November 1990): 973–1002.

³ Stephen A. Woodbury and Robert G. Spiegelman, “Bonuses to Workers and Employers to Reduce Unemployment: Randomized Trials in Illinois,” *American Economic Review* 77 (September 1987): 513–530.

figure 6-3



Real-Wage Rigidity Leads to Job Rationing If the real wage is stuck above the equilibrium level, then the supply of labor exceeds the demand. The result is unemployment.

To understand wage rigidity and structural unemployment, we must examine why the labor market does not clear. When the real wage exceeds the equilibrium level and the supply of workers exceeds the demand, we might expect firms to lower the wages they pay. Structural unemployment arises because firms fail to reduce wages despite an excess supply of labor. We now turn to three causes of this wage rigidity: minimum-wage laws, the monopoly power of unions, and efficiency wages.

Minimum-Wage Laws

The government causes wage rigidity when it prevents wages from falling to equilibrium levels. Minimum-wage laws set a legal minimum on the wages that firms pay their employees. Since the passage of the Fair Labor Standards Act of 1938, the U.S. federal government has enforced a minimum wage that usually has been between 30 and 50 percent of the average wage in manufacturing. For most workers, this minimum wage is not binding, because they earn well above the minimum. Yet for some workers, especially the unskilled and inexperienced, the minimum wage raises their wage above its equilibrium level. It therefore reduces the quantity of their labor that firms demand.

Economists believe that the minimum wage has its greatest impact on teenage unemployment. The equilibrium wages of teenagers tend to be low for two reasons. First, because teenagers are among the least skilled and least experienced members of the labor force, they tend to have low marginal productivity. Second, teenagers often take some of their “compensation” in the

form of on-the-job training rather than direct pay. An apprenticeship is a classic example of training offered in place of wages. For both these reasons, the wage at which the supply of teenage workers equals the demand is low. The minimum wage is therefore more often binding for teenagers than for others in the labor force.

Many economists have studied the impact of the minimum wage on teenage employment. These researchers compare the variation in the minimum wage over time with the variation in the number of teenagers with jobs. These studies find that a 10-percent increase in the minimum wage reduces teenage employment by 1 to 3 percent.⁴

The minimum wage is a perennial source of political debate. Advocates of a higher minimum wage view it as a means of raising the income of the working poor. Certainly, the minimum wage provides only a meager standard of living: in the United States, two adults working full time at minimum-wage jobs would just exceed the official poverty level for a family of four. Although minimum-wage advocates often admit that the policy causes unemployment for some workers, they argue that this cost is worth bearing to raise others out of poverty.

Opponents of a higher minimum wage claim that it is not the best way to help the working poor. They contend not only that the increased labor costs would raise unemployment but also that the minimum wage is poorly targeted. Many minimum-wage earners are teenagers from middle-class homes working for discretionary spending money. Of the approximately 3 million workers who earn the minimum wage, more than one-third are teenagers.

To mitigate the effects on teenage unemployment, some economists and policymakers have long advocated exempting young workers from the regular minimum wage. This would permit a lower wage for teenagers, thereby reducing their unemployment and enabling them to get training and job experience. Opponents of this exemption argue that it gives firms an incentive to substitute teenagers for unskilled adults, thereby raising unemployment among that group. A limited exemption of this kind was tried from 1991 to 1993. Because of many restrictions on its use, however, it had only limited effect and, therefore, was not renewed by Congress.

Many economists and policymakers believe that tax credits are a better way to increase the incomes of the working poor. The *earned income tax credit* is an amount that poor working families are allowed to subtract from the taxes they owe. For a family with a very low income, the credit exceeds its taxes, and the family receives a payment from the government. Unlike the minimum wage, the earned income tax credit does not raise labor costs to firms and, therefore, does not reduce the quantity of labor that firms demand. It has the disadvantage, however, of reducing the government's tax revenue.

⁴ Charles Brown, "Minimum Wage Laws: Are They Overrated?" *Journal of Economic Perspectives* 2 (Summer 1988): 133–146.

CASE STUDY

A Revisionist View of the Minimum Wage

Although most economists believe that increases in the minimum wage reduce employment among workers with little skill and experience, some recent studies question this conclusion. Three respected labor economists—David Card, Lawrence Katz, and Alan Krueger—examined several instances of minimum-wage changes in order to determine the magnitude of the employment response. What they found was startling.

One study examined hiring by fast-food restaurants in New Jersey when New Jersey raised the state minimum wage. Fast-food restaurants are a natural type of firm to examine because they employ many low-wage workers. To control for other effects, such as overall economic conditions, the New Jersey restaurants were compared to similar restaurants across the river in Pennsylvania. Pennsylvania did not raise its minimum wage at the same time. According to standard theory, employment in New Jersey restaurants should have fallen relative to employment in Pennsylvania restaurants. In contrast to this hypothesis, the data showed that employment *rose* in the New Jersey restaurants.

How is this seemingly perverse result possible? One explanation is that firms have some market power in the labor market. As you may have learned in courses in microeconomics, a monopsony firm buys less labor at a lower wage than a competitive firm would. In essence, the firm reduces employment in order to depress the wage it has to pay. A minimum wage prevents the monopsony firm from following this strategy and so (up to a point) can increase employment.

This new view of the minimum wage is controversial. Critics have questioned the reliability of the data used in the New Jersey study. Some studies using other data sources have reached the traditional conclusion that the minimum wage depresses employment. Moreover, most economists are skeptical of the monopsony explanation, because most firms compete with many other firms for workers. Yet this new view has directly affected the policy debate. Lawrence Katz was the first chief economist in the Department of Labor during the Clinton administration. He was followed in this job by Alan Krueger. It is therefore not surprising that President Clinton supported increases in the national minimum wage.⁵

Unions and Collective Bargaining

A second cause of wage rigidity is the monopoly power of unions. Table 6-1 shows the importance of unions in 12 major countries. In the United States, only 16 percent of workers belong to unions. In most European countries, unions play a much larger role.

⁵ To read more about this new view of the minimum wage, see David Card and Alan Krueger, *Myth and Measurement: The New Economics of the Minimum Wage* (Princeton, N.J.: Princeton University Press, 1995); and Lawrence Katz and Alan Krueger, “The Effects of the Minimum Wage on the Fast-Food Industry,” *Industrial and Labor Relations Review* 46 (October 1992): 6–21.

table 6-1

Union Membership as a Percentage of Employment

Country	Percentage Union Workers	Country	Percentage Union Workers
Sweden	84	Germany	33
Denmark	75	Netherlands	28
Italy	47	Switzerland	28
United Kingdom	41	Japan	26
Australia	34	United States	16
Canada	33	France	11

Source: Clara Chang and Constance Sorrentino, "Union Membership Statistics in 12 Countries," *Monthly Labor Review* (December 1991): 46–53.

The wages of unionized workers are determined not by the equilibrium of supply and demand but by collective bargaining between union leaders and firm management. Often, the final agreement raises the wage above the equilibrium level and allows the firm to decide how many workers to employ. The result is a reduction in the number of workers hired, a lower rate of job finding, and an increase in structural unemployment.

Unions can also influence the wages paid by firms whose workforces are not unionized because the threat of unionization can keep wages above the equilibrium level. Most firms dislike unions. Unions not only raise wages but also increase the bargaining power of labor on many other issues, such as hours of employment and working conditions. A firm may choose to pay its workers high wages to keep them happy in order to discourage them from forming a union.

The unemployment caused by unions and by the threat of unionization is an instance of conflict between different groups of workers—**insiders** and **outsiders**. Those workers already employed by a firm, the insiders, typically try to keep their firm's wages high. The unemployed, the outsiders, bear part of the cost of higher wages because at a lower wage they might be hired. These two groups inevitably have conflicting interests. The effect of any bargaining process on wages and employment depends crucially on the relative influence of each group.

The conflict between insiders and outsiders is resolved differently in different countries. In some countries, such as the United States, wage bargaining takes place at the level of the firm or plant. In other countries, such as Sweden, wage bargaining takes place at the national level—with the government often playing a key role. Despite a highly unionized labor force, Sweden has not experienced extraordinarily high unemployment throughout its history. One possible explanation is that the centralization of wage bargaining and the role of the government in the bargaining process give more influence to the outsiders, which keeps wages closer to the equilibrium level.

Efficiency Wages

Efficiency-wage theories propose a third cause of wage rigidity in addition to minimum-wage laws and unionization. These theories hold that high wages make workers more productive. The influence of wages on worker efficiency may explain the failure of firms to cut wages despite an excess supply of labor. Even though a wage reduction would lower a firm's wage bill, it would also—if these theories are correct—lower worker productivity and the firm's profits.

Economists have proposed various theories to explain how wages affect worker productivity. One efficiency-wage theory, which is applied mostly to poorer countries, holds that wages influence nutrition. Better-paid workers can afford a more nutritious diet, and healthier workers are more productive. A firm may decide to pay a wage above the equilibrium level to maintain a healthy workforce. Obviously, this consideration is not important for employers in wealthy countries, such as the United States and most of Europe, because the equilibrium wage is well above the level necessary to maintain good health.

A second efficiency-wage theory, which is more relevant for developed countries, holds that high wages reduce labor turnover. Workers quit jobs for many reasons—to accept better positions at other firms, to change careers, or to move to other parts of the country. The more a firm pays its workers, the greater their incentive to stay with the firm. By paying a high wage, a firm reduces the frequency of quits, thereby decreasing the time spent hiring and training new workers.

A third efficiency-wage theory holds that the average quality of a firm's workforce depends on the wage it pays its employees. If a firm reduces its wage, the best employees may take jobs elsewhere, leaving the firm with inferior employees who have fewer alternative opportunities. Economists recognize this unfavorable sorting as an example of *adverse selection*—the tendency of people with more information (in this case, the workers, who know their own outside opportunities) to self-select in a way that disadvantages people with less information (the firm). By paying a wage above the equilibrium level, the firm may reduce adverse selection, improve the average quality of its workforce, and thereby increase productivity.

A fourth efficiency-wage theory holds that a high wage improves worker effort. This theory posits that firms cannot perfectly monitor their employees' work effort, and that employees must themselves decide how hard to work. Workers can choose to work hard, or they can choose to shirk and risk getting caught and fired. Economists recognize this possibility as an example of *moral hazard*—the tendency of people to behave inappropriately when their behavior is imperfectly monitored. The firm can reduce the problem of moral hazard by paying a high wage. The higher the wage, the greater the cost to the worker of getting fired. By paying a higher wage, a firm induces more of its employees not to shirk and thus increases their productivity.

Although these four efficiency-wage theories differ in detail, they share a common theme: because a firm operates more efficiently if it pays its workers a high wage, the firm may find it profitable to keep wages above the level that balances supply and demand. The result of this higher-than-equilibrium wage is a lower rate of job finding and greater unemployment.⁶

CASE STUDY

Henry Ford's \$5 Workday

In 1914 the Ford Motor Company started paying its workers \$5 per day. The prevailing wage at the time was between \$2 and \$3 per day, so Ford's wage was well above the equilibrium level. Not surprisingly, long lines of job seekers waited outside the Ford plant gates hoping for a chance to earn this high wage.

What was Ford's motive? Henry Ford later wrote, "We wanted to pay these wages so that the business would be on a lasting foundation. We were building for the future. A low wage business is always insecure. . . . The payment of five dollars a day for an eight hour day was one of the finest cost cutting moves we ever made."

From the standpoint of traditional economic theory, Ford's explanation seems peculiar. He was suggesting that *high* wages imply *low* costs. But perhaps Ford had discovered efficiency-wage theory. Perhaps he was using the high wage to increase worker productivity.

Evidence suggests that paying such a high wage did benefit the company. According to an engineering report written at the time, "The Ford high wage does away with all the inertia and living force resistance. . . . The workingmen are absolutely docile, and it is safe to say that since the last day of 1913, every single day has seen major reductions in Ford shops' labor costs." Absenteeism fell by 75 percent, suggesting a large increase in worker effort. Alan Nevins, a historian who studied the early Ford Motor Company, wrote, "Ford and his associates freely declared on many occasions that the high wage policy had turned out to be good business. By this they meant that it had improved the discipline of the workers, given them a more loyal interest in the institution, and raised their personal efficiency."⁷

6-4 Patterns of Unemployment

So far we have developed the theory behind the natural rate of unemployment. We began by showing that the economy's steady-state unemployment rate depends on the rates of job separation and job finding. Then we discussed two reasons why job finding is not instantaneous: the process of job search (which leads to frictional unemployment) and wage rigidity (which leads to structural unemployment). Wage rigidity, in turn, arises from minimum-wage laws, unionization, and efficiency wages.

With these theories as background, we now examine some additional facts about unemployment. These facts will help us evaluate our theories and assess public policies aimed at reducing unemployment.

⁶ For more extended discussions of efficiency wages, see Janet Yellen, "Efficiency Wage Models of Unemployment," *American Economic Review Papers and Proceedings* (May 1984): 200–205; and Lawrence Katz, "Efficiency Wages: A Partial Evaluation," *NBER Macroeconomics Annual* (1986): 235–276.

⁷ Jeremy I. Bulow and Lawrence H. Summers, "A Theory of Dual Labor Markets With Application to Industrial Policy, Discrimination, and Keynesian Unemployment," *Journal of Labor Economics* 4 (July 1986): 376–414; and Daniel M. G. Raff and Lawrence H. Summers, "Did Henry Ford Pay Efficiency Wages?" *Journal of Labor Economics* 5 (October 1987, Part 2): S57–S86.

The Duration of Unemployment

When a person becomes unemployed, is the spell of unemployment likely to be short or long? The answer to this question is important because it indicates the reasons for the unemployment and what policy response is appropriate. On the one hand, if most unemployment is short term, one might argue that it is frictional and perhaps unavoidable. Unemployed workers may need some time to search for the job that is best suited to their skills and tastes. On the other hand, long-term unemployment cannot easily be attributed to the time it takes to match jobs and workers: we would not expect this matching process to take many months. Long-term unemployment is more likely to be structural unemployment. Thus, data on the duration of unemployment can affect our view about the reasons for unemployment.

The answer to our question turns out to be subtle. The data show that most spells of unemployment are short but that most weeks of unemployment are attributable to the long-term unemployed. Consider the data for a typical year, 1974, during which the unemployment rate was 5.6 percent. In that year, 60 percent of the spells of unemployment ended within one month, yet 69 percent of the weeks of unemployment occurred in spells that lasted two or more months.⁸

To see how both these facts can be true, consider the following example. Suppose that 10 people are unemployed for part of a given year. Of these 10 people, 8 are unemployed for 1 month, and 2 are unemployed for 12 months, totaling 32 months of unemployment. In this example, most spells of unemployment are short: 8 of the 10 unemployment spells, or 80 percent, end in 1 month. Yet most months of unemployment are attributable to the long-term unemployed: 24 of the 32 months of unemployment, or 75 percent, are experienced by the 2 workers who are unemployed for 12 months. Depending on whether we look at spells of unemployment or months of unemployment, most unemployment can appear to be short term or long term.

This evidence on the duration of unemployment has an important implication for public policy. If the goal is to lower substantially the natural rate of unemployment, policies must aim at the long-term unemployed, because these individuals account for a large amount of unemployment. Yet policies must be carefully targeted, because the long-term unemployed constitute a small minority of those who become unemployed. Most people who become unemployed find work within a short time.

Variation in the Unemployment Rate Across Demographic Groups

The rate of unemployment varies substantially across different groups within the population. Table 6-2 presents the U.S. unemployment rates for different demographic groups in 2000, when the overall unemployment rate was 4.0 percent.

This table shows that younger workers have much higher unemployment rates than older ones. To explain this difference, recall our model of the natural rate of

⁸ Kim B. Clark and Lawrence H. Summers, "Labor Market Dynamics and Unemployment: A Reconsideration," *Brookings Papers on Economic Activity* (1979:1): 13–72.

table 6-2

Unemployment Rate by Demographic Group: 2000

Age	White Male	White Female	Black Male	Black Female
16–19	12.3	10.4	26.4	23.0
20 and over	2.8	3.1	7.0	6.3

Source: U.S. Department of Labor.

unemployment. The model isolates two possible causes for a high rate of unemployment: a low rate of job finding and a high rate of job separation. When economists study data on the transition of individuals between employment and unemployment, they find that those groups with high unemployment tend to have high rates of job separation. They find less variation across groups in the rate of job finding. For example, employed white males are four times more likely to become unemployed if they are teenagers than if they are middle-aged; once someone is unemployed, the rate of job finding is not closely related to age.

These findings help explain the higher unemployment rates for younger workers. Younger workers have only recently entered the labor market, and they are often uncertain about their career plans. It may be best for them to try different types of jobs before making a long-term commitment to a specific occupation. If so, we should expect a higher rate of job separation and a higher rate of frictional unemployment for this group.

Another fact that stands out from Table 6-2 is that unemployment rates are much higher for blacks than for whites. This phenomenon is not well understood. Data on transitions between employment and unemployment show that the higher unemployment rates for blacks, and especially for black teenagers, arise because of both higher rates of job separation and lower rates of job finding. Possible reasons for the lower rates of job finding include less access to informal job-finding networks and discrimination by employers.

Trends in U.S. Unemployment

Over the past half century, the natural rate of unemployment in the United States has not been stable. If you look back at Figure 6-1, you will see that unemployment averaged below 5 percent in the 1950s and 1960s, rose to over 6 percent in the 1970s and 1980s, and then drifted back below 5 percent in the 1990s. Although economists do not have a conclusive explanation for these changes, they have proposed several hypotheses.

Demographics One explanation stresses the changing composition of the U.S. labor force. After World War II, birthrates rose dramatically: the number of births rose from 2.9 million in 1945 to a peak of 4.3 million in 1957, before falling back to 3.1 million in 1973. This rise in births in the 1950s led to a rise in the number of young workers in the 1970s. Younger workers have higher unemployment rates, however, so when the baby-boom generation entered the labor force,

they increased the average level of unemployment. Then as the baby-boom workers aged, the average age of the labor force increased, lowering the average unemployment rate in the 1990s.

This demographic change, however, cannot fully explain the trends in unemployment because similar trends are apparent for fixed demographic groups. For example, for men between the ages of 25 and 54, the average unemployment rate rose from 3.0 percent in the 1960s to 6.1 percent in the 1980s. Thus, although demographic changes may be part of the story of rising unemployment over this period, there must be other explanations of the long-term trend as well.

Sectoral Shifts A second explanation is based on changes in the prevalence of sectoral shifts. The greater the amount of sectoral reallocation, the greater the rate of job separation and the higher the level of frictional unemployment. One source of sectoral shifts during the 1970s and early 1980s was the great volatility in oil prices caused by OPEC, the international oil cartel. These large changes in oil prices may have required reallocating labor between more-energy-intensive and less-energy-intensive sectors. If so, oil-price volatility may have increased unemployment during this period. Although this explanation is hard to evaluate, it is consistent with recent developments: the fall in unemployment during the 1990s coincided with increased stability in oil prices.

Productivity A third explanation for the trends in unemployment emphasizes the link between unemployment and productivity. As Chapter 8 discusses more fully, the 1970s experienced a slowdown in productivity growth, and the 1990s experienced a pickup in productivity growth. These productivity changes roughly coincide with changes in unemployment. Perhaps slowing productivity during the 1970s raised the natural rate of unemployment, and accelerating productivity during the 1990s growth lowered it.

Why such an effect would occur, however, is not obvious. In standard theories of the labor market, higher productivity means greater labor demand and thus higher real wages, but unemployment is unchanged. This prediction is consistent with the long-term data, which show consistent upward trends in productivity and real wages but no trend in unemployment. Yet suppose that workers are slow to catch on to news about productivity. When productivity changes, workers may only gradually alter the real wages they ask from their employers, making real wages sluggish in response to labor demand. An acceleration in productivity growth, such as that experienced during the 1990s, will increase labor demand and, with a sluggish real wage, reduce the amount of unemployment.

In the end, the trends in the unemployment rate remain a mystery. The proposed explanations are plausible, but none seems conclusive on its own. Perhaps there is no single answer. The upward drift in the unemployment rate in the 1970s and 1980s and the downward drift in the 1990s may be the result of several unrelated developments.⁹

⁹ On the role of demographics, see Robert Shimer, “Why Is the U.S. Unemployment Rate So Much Lower?” *NBER Macroeconomics Annual* 13 (1998). On the role of sectoral shifts, see David M. Lilien, “Sectoral Shifts and Cyclical Unemployment,” *Journal of Political Economy* 90 (August 1982): 777–793. On the role of productivity, see Laurence Ball and Robert Moffitt, “Productivity Growth and the Phillips Curve,” *NBER Working Paper No. 8421*, August 2001.

Transitions Into and Out of the Labor Force

So far we have ignored an important aspect of labor-market dynamics: the movement of individuals into and out of the labor force. Our model of the natural rate of unemployment assumes that the size of the labor force is fixed. In this case, the sole reason for unemployment is job separation, and the sole reason for leaving unemployment is job finding.

In fact, changes in the labor force are important. About one-third of the unemployed have only recently entered the labor force. Some of these entrants are young workers still looking for their first jobs; others have worked before but temporarily left the labor force. In addition, not all unemployment ends with job finding; almost half of all spells of unemployment end in the unemployed person's withdrawal from the labor market.

Individuals entering and leaving the labor force make unemployment statistics more difficult to interpret. On the one hand, some individuals calling themselves unemployed may not be seriously looking for jobs and perhaps should best be viewed as out of the labor force. Their "unemployment" may not represent a social problem. On the other hand, some individuals may want jobs but, after unsuccessful searches, have given up looking. These **discouraged workers** are counted as being out of the labor force and do not show up in unemployment statistics. Even though their joblessness is unmeasured, it may nonetheless be a social problem.

Because of these and many other issues that complicate the interpretation of the unemployment data, the Bureau of Labor Statistics calculates several measures of labor underutilization. Table 6-3 gives the definitions and their values as

table 6-3

Alternative Measures of Labor Underutilization

Definition		Percentage in March 2001
U-1	Persons unemployed 15 weeks or longer, as a percentage of the civilian labor force (includes only very long term unemployed)	1.2 %
U-2	Job losers and persons who have completed temporary jobs, as a percentage of the civilian labor force (excludes job leavers)	2.4
U-3	Total unemployed, as a percentage of the civilian labor force (official unemployment rate)	4.6
U-4	Total unemployed, plus discouraged workers, as a percentage of the civilian labor force plus discouraged workers	4.8
U-5	Total unemployed plus all marginally attached workers, as a percentage of the civilian labor force plus all marginally attached workers	5.3
U-6	Total unemployed, plus all marginally attached workers, plus total employed part time for economic reasons, as a percentage of the civilian labor force plus all marginally attached workers	7.6

Note: *Marginally attached workers* are persons who currently are neither working nor looking for work but indicate that they want and are available for a job and have looked for work sometime in the recent past. *Discouraged workers*, a subset of the marginally attached, have given a job-market related reason for not currently looking for a job. *Persons employed part time for economic reasons* are those who want and are available for full-time work but have had to settle for a part-time schedule.

Source: U.S. Department of Labor.

of March 2001. The measures range from 1.2 to 7.6 percent, depending on the characteristics one uses to classify a worker as not fully employed.


The Rise in European Unemployment

Although our discussion has focused largely on the United States, one puzzling question about unemployment concerns recent developments in Europe. Figure 6-4 shows the rate of unemployment in the countries that make up the European Community—Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Portugal, Spain, Sweden, and the United Kingdom. As you can see, the rate of unemployment in these countries has risen substantially: it averaged less than 3 percent in the 1960s and more than 10 percent in recent years.

What is the cause of rising European unemployment? No one knows for sure, but there is a leading theory. Many economists believe that the problem can be traced to generous benefits for unemployed workers, coupled with a technologically driven fall in the demand for unskilled workers relative to skilled workers.

There is no question that most European countries have generous programs for those without jobs. These programs go by various names: social insurance, the welfare state, or simply “the dole.” Many countries allow the unemployed to

figure 6-4



Unemployment in the European Community This figure shows the unemployment rate in the 15 countries that make up the European Community. The figure shows that the European unemployment rate has risen substantially since 1980.

Source: OECD.

collect benefits indefinitely, rather than for only a short period of time as in the United States. Studies have shown that countries with more generous benefits tend to have higher rates of unemployment. In some sense, those living on the dole are really out of the labor force: given the employment opportunities available, taking a job is less attractive than remaining without work. Yet these people are often counted as unemployed in government statistics.

There is also no question that the demand for unskilled workers has fallen relative to the demand for skilled workers. This change in demand is probably attributable to changes in technology: computers, for example, increase the demand for workers who can use them and reduce the demand for those who cannot. In the United States, this change in demand has been reflected in wages rather than unemployment: over the past two decades, the wages of unskilled workers have fallen substantially relative to the wages of skilled workers. In Europe, however, the welfare state provides unskilled workers with an alternative to working for low wages. As the wages of unskilled workers fall, more workers view the dole as their best available option. The result is higher unemployment.

This diagnosis of high European unemployment does not suggest an easy remedy. Reducing the magnitude of government benefits for the unemployed would encourage workers to get off the dole and accept low-wage jobs. But it would also exacerbate economic inequality—the very problem that welfare-state policies were designed to address.¹⁰

CASE STUDY

The Secrets to Happiness

Why are some people more satisfied with their lives than others? This is a deep and difficult question, most often left to philosophers and psychologists. But part of the answer is macroeconomic. Recent research has shown that people are happier when they are living in a country with low inflation and low unemployment.

From 1975 to 1991, a survey called the Euro-Barometer Survey Series asked 264,710 people living in 12 European countries about their happiness and overall satisfaction with life. One question asked, “On the whole, are you very satisfied, fairly satisfied, not very satisfied, or not at all satisfied with the life you lead?” To see what determines happiness, the answers to this question were correlated with individual and macroeconomic variables. Other things being equal, people are more satisfied with their lives if they are rich, educated, married, in school, self-employed, retired, female, and young or old (as opposed to middle-aged).

¹⁰ For more discussion of these issues, see Paul Krugman, “Past and Prospective Causes of High Unemployment,” in *Reducing Unemployment: Current Issues and Policy Options*, Federal Reserve Bank of Kansas City, August 1994.

They are less satisfied if they are unemployed, divorced, or living with adolescent children. (Some of these correlations may reflect the effects, rather than causes, of happiness: for example, a happy person may find it easier than an unhappy one to keep a job and a spouse.)

Beyond these individual characteristics, the economy's overall rates of unemployment and inflation also play a significant role in explaining reported happiness. An increase in the unemployment rate of 4 percentage points is large enough to move 11 percent of the population down from one life-satisfaction category to another. The overall unemployment rate reduces satisfaction even after controlling for an individual's employment status. That is, the employed in a high-unemployment nation are less happy than their counterparts in a low-unemployment nation, perhaps because they are more worried about job loss or perhaps out of sympathy with their fellow citizens.

High inflation is also associated with lower life satisfaction, although the effect is not as large. A 1.7-percentage-point increase in inflation reduces happiness by about as much as a 1-percentage-point increase in unemployment. The commonly cited "misery index," which is the sum of the inflation and unemployment rates, apparently gives too much weight to inflation relative to unemployment.¹¹

6-5 Conclusion

Unemployment represents wasted resources. Unemployed workers have the potential to contribute to national income but are not doing so. Those searching for jobs to suit their skills are happy when the search is over, and those waiting for jobs in firms that pay above-equilibrium wages are happy when positions open up.

Unfortunately, neither frictional unemployment nor structural unemployment can be easily reduced. The government cannot make job search instantaneous, and it cannot easily bring wages closer to equilibrium levels. Zero unemployment is not a plausible goal for free-market economies.

Yet public policy is not powerless in the fight to reduce unemployment. Job-training programs, the unemployment-insurance system, the minimum wage, and the laws governing collective bargaining are often topics of political debate. The policies we choose are likely to have important effects on the economy's natural rate of unemployment.

¹¹ Rafael Di Tella, Robert J. MacCulloch, and Andrew J. Oswald, "Preferences Over Inflation and Unemployment: Evidence From Surveys of Happiness," *American Economic Review* 91 (March 2001): 335–341.

Summary

1. The natural rate of unemployment is the steady-state rate of unemployment. It depends on the rate of job separation and the rate of job finding.
2. Because it takes time for workers to search for the job that best suits their individual skills and tastes, some frictional unemployment is inevitable. Various government policies, such as unemployment insurance, alter the amount of frictional unemployment.
3. Structural unemployment results when the real wage remains above the level that equilibrates labor supply and labor demand. Minimum-wage legislation is one cause of wage rigidity. Unions and the threat of unionization are another. Finally, efficiency-wage theories suggest that, for various reasons, firms may find it profitable to keep wages high despite an excess supply of labor.
4. Whether we conclude that most unemployment is short term or long term depends on how we look at the data. Most spells of unemployment are short. Yet most weeks of unemployment are attributable to the small number of long-term unemployed.
5. The unemployment rates among demographic groups differ substantially. In particular, the unemployment rates for younger workers are much higher than for older workers. This results from a difference in the rate of job separation rather than from a difference in the rate of job finding.
6. The natural rate of unemployment in the United States has exhibited long-term trends. In particular, it rose from the 1950s to the 1970s and then started drifting downward again in the 1990s. Various explanations have been proposed, including the changing demographic composition of the labor force, changes in the prevalence of sectoral shifts, and changes in the rate of productivity growth.
7. Individuals who have recently entered the labor force, including both new entrants and reentrants, make up about one-third of the unemployed. Transitions into and out of the labor force make unemployment statistics more difficult to interpret.

KEY CONCEPTS

Natural rate of unemployment

Frictional unemployment

Sectoral shift

Unemployment insurance

Wage rigidity

Structural unemployment

Insiders versus outsiders

Efficiency wages

Discouraged workers

QUESTIONS FOR REVIEW

1. What determines the natural rate of unemployment?
2. Describe the difference between frictional unemployment and structural unemployment.
3. Give three explanations why the real wage may remain above the level that equilibrates labor supply and labor demand.
4. Is most unemployment long term or short term? Explain your answer.
5. How do economists explain the high natural rate of unemployment in the 1970s and 1980s? How do they explain the fall in the natural rate in the 1990s?

PROBLEMS AND APPLICATIONS

1. Answer the following questions about your own experience in the labor force:
 - a. When you or one of your friends is looking for a part-time job, how many weeks does it typically take? After you find a job, how many weeks does it typically last?
 - b. From your estimates, calculate (in a rate per week) your rate of job finding f and your rate of job separation s . (*Hint:* If f is the rate of job finding, then the average spell of unemployment is $1/f$.)
 - c. What is the natural rate of unemployment for the population you represent?
2. In this chapter we saw that the steady-state rate of unemployment is $U/L = s/(s + f)$. Suppose that the unemployment rate does not begin at this level. Show that unemployment will evolve over time and reach this steady state. (*Hint:* Express the change in the number of unemployed as a function of s , f , and U . Then show that if unemployment is above the natural rate, unemployment falls, and if unemployment is below the natural rate, unemployment rises.)
3. The residents of a certain dormitory have collected the following data: People who live in the dorm can be classified as either involved in a relationship or uninvolved. Among involved people, 10 percent experience a breakup of their relationship every month. Among uninvolved people, 5 percent will enter into a relationship every month. What is the steady-state fraction of residents who are uninvolved?
4. Suppose that Congress passes legislation making it more difficult for firms to fire workers. (An example is a law requiring severance pay for fired workers.) If this legislation reduces the rate of job separation without affecting the rate of job finding, how would the natural rate of unemployment change? Do you think that it is plausible that the legislation would not affect the rate of job finding? Why or why not?
5. Consider an economy with the following Cobb–Douglas production function:

$$Y = K^{1/3}L^{2/3}.$$

The economy has 1,000 units of capital and a labor force of 1,000 workers.

- a. Derive the equation describing labor demand in this economy as a function of the real wage and the capital stock. (*Hint:* Review the appendix to Chapter 3.)
- b. If the real wage can adjust to equilibrate labor supply and labor demand, what is the real wage? In this equilibrium, what are employment, output, and the total amount earned by workers?
- c. Now suppose that Congress, concerned about the welfare of the working class, passes a law requiring firms to pay workers a real wage of 1 unit of output. How does this wage compare to the equilibrium wage?
- d. Congress cannot dictate how many workers firms hire at the mandated wage. Given this fact, what are the effects of this law? Specifi-

- cally, what happens to employment, output, and the total amount earned by workers?
- e. Will Congress succeed in its goal of helping the working class? Explain.
 - f. Do you think that this analysis provides a good way of thinking about a minimum-wage law? Why or why not?
6. Suppose that a country experiences a reduction in productivity—that is, an adverse shock to the production function.
- a. What happens to the labor demand curve?
 - b. How would this change in productivity affect the labor market—that is, employment, unemployment, and real wages—if the labor market were always in equilibrium?
 - c. How would this change in productivity affect the labor market if unions prevented real wages from falling?
7. In any city at any time, some of the stock of usable office space is vacant. This vacant office space is unemployed capital. How would you explain this phenomenon? Is it a social problem?

part III

Growth Theory: The Economy in the Very Long Run

Economic Growth I

The question of growth is nothing new but a new disguise for an age-old issue, one which has always intrigued and preoccupied economics: the present versus the future.

—James Tobin

If you have ever spoken with your grandparents about what their lives were like when they were young, most likely you learned an important lesson about economics: material standards of living have improved substantially over time for most families in most countries. This advance comes from rising incomes, which have allowed people to consume greater quantities of goods and services.

To measure economic growth, economists use data on gross domestic product, which measures the total income of everyone in the economy. The real GDP of the United States today is more than three times its 1950 level, and real GDP per person is more than twice its 1950 level. In any given year, we can also observe large differences in the standard of living among countries. Table 7-1 shows income per person in 1999 of the world's 12 most populous countries. The United States tops the list with an income of \$31,910 per person. Nigeria has an income per person of only \$770—less than 3 percent of the figure for the United States.

Our goal in this part of the book is to understand what causes these differences in income over time and across countries. In Chapter 3 we identified the factors of production—capital and labor—and the production technology as the sources of the economy's output and, thus, of its total income. Differences in income, then, must come from differences in capital, labor, and technology.

Our primary task is to develop a theory of economic growth called the **Solow growth model**. Our analysis in Chapter 3 enabled us to describe how the economy produces and uses its output at one point in time. The analysis was static—a snapshot of the economy. To explain why our national income grows, and why some economies grow faster than others, we must broaden our analysis so that it describes changes in the economy over time. By developing such a model, we make our analysis dynamic—more like a movie than a photograph. The Solow growth model shows how saving, population growth, and

table 7-1**International Differences in the Standard of Living: 1999**

Country	Income per Person (in U.S. dollars)	Country	Income per Person (in U.S. dollars)
United States	\$31,910	China	3,550
Japan	25,170	Indonesia	2,660
Germany	23,510	India	2,230
Mexico	8,070	Pakistan	1,860
Russia	6,990	Bangladesh	1,530
Brazil	6,840	Nigeria	770

Source: World Bank.

technological progress affect the level of an economy's output and its growth over time. In this chapter we analyze the roles of saving and population growth. In the next chapter we introduce technological progress.¹

7-1 The Accumulation of Capital

The Solow growth model is designed to show how growth in the capital stock, growth in the labor force, and advances in technology interact in an economy, and how they affect a nation's total output of goods and services. We build this model in steps. Our first step is to examine how the supply and demand for goods determine the accumulation of capital. In this first step, we assume that the labor force and technology are fixed. We then relax these assumptions by introducing changes in the labor force later in this chapter and by introducing changes in technology in the next.

The Supply and Demand for Goods

The supply and demand for goods played a central role in our static model of the closed economy in Chapter 3. The same is true for the Solow model. By considering the supply and demand for goods, we can see what determines how much output is produced at any given time and how this output is allocated among alternative uses.

¹ The Solow growth model is named after economist Robert Solow and was developed in the 1950s and 1960s. In 1987 Solow won the Nobel Prize in economics for his work in economic growth. The model was introduced in Robert M. Solow, "A Contribution to the Theory of Economic Growth," *Quarterly Journal of Economics* (February 1956): 65–94.

The Supply of Goods and the Production Function The supply of goods in the Solow model is based on the now-familiar production function, which states that output depends on the capital stock and the labor force:

$$Y = F(K, L).$$

The Solow growth model assumes that the production function has constant returns to scale. This assumption is often considered realistic, and as we will see shortly, it helps simplify the analysis. Recall that a production function has constant returns to scale if

$$zY = F(zK, zL)$$

for any positive number z . That is, if we multiply both capital and labor by z , we also multiply the amount of output by z .

Production functions with constant returns to scale allow us to analyze all quantities in the economy relative to the size of the labor force. To see that this is true, set $z = 1/L$ in the preceding equation to obtain

$$Y/L = F(K/L, 1).$$

This equation shows that the amount of output per worker Y/L is a function of the amount of capital per worker K/L . (The number “1” is, of course, constant and thus can be ignored.) The assumption of constant returns to scale implies that the size of the economy—as measured by the number of workers—does not affect the relationship between output per worker and capital per worker.

Because the size of the economy does not matter, it will prove convenient to denote all quantities in per-worker terms. We designate these with lowercase letters, so $y = Y/L$ is output per worker, and $k = K/L$ is capital per worker. We can then write the production function as

$$y = f(k),$$

where we define $f(k) = F(k, 1)$. Figure 7-1 illustrates this production function.


The slope of this production function shows how much extra output a worker produces when given an extra unit of capital. This amount is the marginal product of capital MPK . Mathematically, we write

$$MPK = f(k + 1) - f(k).$$

Note that in Figure 7-1, as the amount of capital increases, the production function becomes flatter, indicating that the production function exhibits diminishing marginal product of capital. When k is low, the average worker has only a little capital to work with, so an extra unit of capital is very useful and produces a lot of additional output. When k is high, the average worker has a lot of capital, so an extra unit increases production only slightly.

The Demand for Goods and the Consumption Function The demand for goods in the Solow model comes from consumption and investment. In other

figure 7-1



The Production Function The production function shows how the amount of capital per worker k determines the amount of output per worker $y = f(k)$. The slope of the production function is the marginal product of capital: if k increases by 1 unit, y increases by MPK units. The production function becomes flatter as k increases, indicating diminishing marginal product of capital.

words, output per worker y is divided between consumption per worker c and investment per worker i :

$$y = c + i.$$

This equation is the per-worker version of the national income accounts identity for an economy. Notice that it omits government purchases (which for present purposes we can ignore) and net exports (because we are assuming a closed economy).

The Solow model assumes that each year people save a fraction s of their income and consume a fraction $(1 - s)$. We can express this idea with a consumption function with the simple form

$$c = (1 - s)y,$$

where s , the saving rate, is a number between zero and one. Keep in mind that various government policies can potentially influence a nation's saving rate, so one of our goals is to find what saving rate is desirable. For now, however, we just take the saving rate s as given.

To see what this consumption function implies for investment, substitute $(1 - s)y$ for c in the national income accounts identity:

$$y = (1 - s)y + i.$$

Rearrange the terms to obtain

$$i = sy.$$

This equation shows that investment equals saving, as we first saw in Chapter 3. Thus, the rate of saving s is also the fraction of output devoted to investment.

We have now introduced the two main ingredients of the Solow model—the production function and the consumption function—which describe the economy at any moment in time. For any given capital stock k , the production function $y = f(k)$ determines how much output the economy produces, and the saving rate s determines the allocation of that output between consumption and investment.

Growth in the Capital Stock and the Steady State


At any moment, the capital stock is a key determinant of the economy's output, but the capital stock can change over time, and those changes can lead to economic growth. In particular, two forces influence the capital stock: investment and depreciation. *Investment* refers to the expenditure on new plant and equipment, and it causes the capital stock to rise. *Depreciation* refers to the wearing out of old capital, and it causes the capital stock to fall. Let's consider each of these in turn.

As we have already noted, investment per worker i equals sy . By substituting the production function for y , we can express investment per worker as a function of the capital stock per worker:

$$i = sf(k).$$


This equation relates the existing stock of capital k to the accumulation of new capital i . Figure 7-2 shows this relationship. This figure illustrates how, for any value of k , the amount of output is determined by the production function $f(k)$, and the allocation of that output between consumption and saving is determined by the saving rate s .

figure 7-2



Output, Consumption, and Investment The saving rate s determines the allocation of output between consumption and investment. For any level of capital k , output is $f(k)$, investment is $sf(k)$, and consumption is $f(k) - sf(k)$.

figure 7-3



Depreciation A constant fraction δ of the capital stock wears out every year. Depreciation is therefore proportional to the capital stock.

To incorporate depreciation into the model, we assume that a certain fraction δ of the capital stock wears out each year. Here δ (the lowercase Greek letter delta) is called the *depreciation rate*. For example, if capital lasts an average of 25 years, then the depreciation rate is 4 percent per year ($\delta = 0.04$). The amount of capital that depreciates each year is δk . Figure 7-3 shows how the amount of depreciation depends on the capital stock.

We can express the impact of investment and depreciation on the capital stock with this equation:

$$\begin{array}{ccccccc} \text{Change in Capital Stock} & = & \text{Investment} & - & \text{Depreciation} \\ \Delta k & = & i & - & \delta k, \end{array}$$

where Δk is the change in the capital stock between one year and the next. Because investment i equals $sf(k)$, we can write this as

$$\Delta k = sf(k) - \delta k.$$


Figure 7-4 graphs the terms of this equation—investment and depreciation—for different levels of the capital stock k . The higher the capital stock, the greater the amounts of output and investment. Yet the higher the capital stock, the greater also the amount of depreciation.

As Figure 7-4 shows, there is a single capital stock k^* at which the amount of investment equals the amount of depreciation. If the economy ever finds itself at this level of the capital stock, the capital stock will not change because the two forces acting on it—investment and depreciation—just balance. That is, at k^* , $\Delta k = 0$, so the capital stock k and output $f(k)$ are steady over time (rather than growing or shrinking). We therefore call k^* the **steady-state** level of capital.

The steady state is significant for two reasons. As we have just seen, an economy at the steady state will stay there. In addition, and just as important, an economy not at the steady state will go there. That is, regardless of the level of capital

figure 7-4

Investment and depreciation



Capital stock increases because investment exceeds depreciation.

Steady-state level of capital per worker

Capital stock decreases because depreciation exceeds investment.

Investment, Depreciation, and the Steady State The steady-state level of capital k^* is the level at which investment equals depreciation, indicating that the amount of capital will not change over time. Below k^* , investment exceeds depreciation, so the capital stock grows. Above k^* , investment is less than depreciation, so the capital stock shrinks.

with which the economy begins, it ends up with the steady-state level of capital. In this sense, *the steady state represents the long-run equilibrium of the economy.*

To see why an economy always ends up at the steady state, suppose that the economy starts with less than the steady-state level of capital, such as level k_1 in Figure 7-4. In this case, the level of investment exceeds the amount of depreciation. Over time, the capital stock will rise and will continue to rise—along with output $f(k)$ —until it approaches the steady state k^* .

Similarly, suppose that the economy starts with more than the steady-state level of capital, such as level k_2 . In this case, investment is less than depreciation: capital is wearing out faster than it is being replaced. The capital stock will fall, again approaching the steady-state level. Once the capital stock reaches the steady state, investment equals depreciation, and there is no pressure for the capital stock to either increase or decrease.

Approaching the Steady State: A Numerical Example

Let's use a numerical example to see how the Solow model works and how the economy approaches the steady state. For this example, we assume that the production function is²

$$Y = K^{1/2}L^{1/2}.$$

² If you read the appendix to Chapter 3, you will recognize this as the Cobb–Douglas production function with the parameter α equal to $1/2$.

To derive the per-worker production function $f(k)$, divide both sides of the production function by the labor force L :

$$\frac{Y}{L} = \frac{K^{1/2}L^{1/2}}{L}.$$

Rearrange to obtain

$$\frac{Y}{L} = \left(\frac{K}{L}\right)^{1/2}.$$

Because $y = Y/L$ and $k = K/L$, this becomes

$$y = k^{1/2}.$$

This equation can also be written as

$$y = \sqrt{k}.$$

This form of the production function states that output per worker is equal to the square root of the amount of capital per worker.

To complete the example, let's assume that 30 percent of output is saved ($s = 0.3$), that 10 percent of the capital stock depreciates every year ($\delta = 0.1$), and that the economy starts off with 4 units of capital per worker ($k = 4$). Given these numbers, we can now examine what happens to this economy over time.

We begin by looking at the production and allocation of output in the first year. According to the production function, the 4 units of capital per worker produce 2 units of output per worker. Because 30 percent of output is saved and invested and 70 percent is consumed, $i = 0.6$ and $c = 1.4$. Also, because 10 percent of the capital stock depreciates, $\delta k = 0.4$. With investment of 0.6 and depreciation of 0.4, the change in the capital stock is $\Delta k = 0.2$. The second year begins with 4.2 units of capital per worker.

Table 7-2 shows how the economy progresses year by year. Every year, new capital is added and output grows. Over many years, the economy approaches a steady state with 9 units of capital per worker. In this steady state, investment of 0.9 exactly offsets depreciation of 0.9, so that the capital stock and output are no longer growing.

Following the progress of the economy for many years is one way to find the steady-state capital stock, but there is another way that requires fewer calculations. Recall that

$$\Delta k = sf(k) - \delta k.$$

This equation shows how k evolves over time. Because the steady state is (by definition) the value of k at which $\Delta k = 0$, we know that

$$0 = sf(k^*) - \delta k^*,$$

or, equivalently,

$$\frac{k^*}{f(k^*)} = \frac{s}{\delta}.$$

table 7-2

Approaching the Steady State: A Numerical Example

Assumptions: $y = \sqrt{k}$; $s = 0.3$; $\delta = 0.1$; initial $k = 4.0$

Year	k	y	c	i	δk	Δk
1	4.000	2.000	1.400	0.600	0.400	0.200
2	4.200	2.049	1.435	0.615	0.420	0.195
3	4.395	2.096	1.467	0.629	0.440	0.189
4	4.584	2.141	1.499	0.642	0.458	0.184
5	4.768	2.184	1.529	0.655	0.477	0.178
.						
.						
.						
10	5.602	2.367	1.657	0.710	0.560	0.150
.						
.						
.						
25	7.321	2.706	1.894	0.812	0.732	0.080
.						
.						
.						
100	8.962	2.994	2.096	0.898	0.896	0.002
.						
.						
.						
∞	9.000	3.000	2.100	0.900	0.900	0.000

This equation provides a way of finding the steady-state level of capital per worker, k^* . Substituting in the numbers and production function from our example, we obtain

$$\frac{k^*}{\sqrt{k^*}} = \frac{0.3}{0.1}.$$

Now square both sides of this equation to find

$$k^* = 9.$$

The steady-state capital stock is 9 units per worker. This result confirms the calculation of the steady state in Table 7-2.

CASE STUDY

The Miracle of Japanese and German Growth

Japan and Germany are two success stories of economic growth. Although today they are economic superpowers, in 1945 the economies of both countries were in shambles. World War II had destroyed much of their capital stocks. In the

decades after the war, however, these two countries experienced some of the most rapid growth rates on record. Between 1948 and 1972, output per person grew at 8.2 percent per year in Japan and 5.7 percent per year in Germany, compared to only 2.2 percent per year in the United States.

Are the postwar experiences of Japan and Germany so surprising from the standpoint of the Solow growth model? Consider an economy in steady state. Now suppose that a war destroys some of the capital stock. (That is, suppose the capital stock drops from k^* to k_1 in Figure 7-4.) Not surprisingly, the level of output immediately falls. But if the saving rate—the fraction of output devoted to saving and investment—is unchanged, the economy will then experience a period of high growth. Output grows because, at the lower capital stock, more capital is added by investment than is removed by depreciation. This high growth continues until the economy approaches its former steady state. Hence, although destroying part of the capital stock immediately reduces output, it is followed by higher-than-normal growth. The “miracle” of rapid growth in Japan and Germany, as it is often described in the business press, is what the Solow model predicts for countries in which war has greatly reduced the capital stock.

How Saving Affects Growth

The explanation of Japanese and German growth after World War II is not quite as simple as suggested in the preceding case study. Another relevant fact is that both Japan and Germany save and invest a higher fraction of their output than the United States. To understand more fully the international differences in economic performance, we must consider the effects of different saving rates.


Consider what happens to an economy when its saving rate increases. Figure 7-5 shows such a change. The economy is assumed to begin in a steady state with saving rate s_1 and capital stock k_1^* . When the saving rate increases from s_1 to s_2 , the $sf(k)$ curve shifts upward. At the initial saving rate s_1 and the initial capital stock k_1^* , the amount of investment just offsets the amount of depreciation. Immediately after the saving rate rises, investment is higher, but the capital stock and depreciation are unchanged. Therefore, investment exceeds depreciation. The capital stock will gradually rise until the economy reaches the new steady state k_2^* , which has a higher capital stock and a higher level of output than the old steady state.

The Solow model shows that the saving rate is a key determinant of the steady-state capital stock. *If the saving rate is high, the economy will have a large capital stock and a high level of output. If the saving rate is low, the economy will have a small capital stock and a low level of output.* This conclusion sheds light on many discussions of fiscal policy. As we saw in Chapter 3, a government budget deficit can reduce national saving and crowd out investment. Now we can see that the long-run consequences of a reduced saving rate are a lower capital stock and lower national income. This is why many economists are critical of persistent budget deficits.

What does the Solow model say about the relationship between saving and economic growth? Higher saving leads to faster growth in the Solow model, but

figure 7-5

Investment and depreciation



An Increase in the Saving Rate An increase in the saving rate s implies that the amount of investment for any given capital stock is higher. It therefore shifts the saving function upward. At the initial steady state k_1^* , investment now exceeds depreciation. The capital stock rises until the economy reaches a new steady state k_2^* with more capital and output.

only temporarily. An increase in the rate of saving raises growth only until the economy reaches the new steady state. If the economy maintains a high saving rate, it will maintain a large capital stock and a high level of output, but it will not maintain a high rate of growth forever.

Now that we understand how saving affects growth, we can more fully explain the impressive economic performances of Germany and Japan after World War II. Not only were their initial capital stocks low because of the war, but their steady-state capital stocks were high because of their high saving rates. Both of these facts help explain the rapid growth of these two countries in the 1950s and 1960s.


CASE STUDY

Saving and Investment Around the World

We started this chapter with an important question: Why are some countries so rich while others are mired in poverty? Our analysis has taken us a step closer to the answer. According to the Solow model, if a nation devotes a large fraction of its income to saving and investment, it will have a high steady-state capital stock and a high level of income. If a nation saves and invests only a small fraction of its income, its steady-state capital and income will be low.

Let's now look at some data to see if this theoretical result in fact helps explain the large international variation in standards of living. Figure 7-6 is a scatterplot of data from 84 countries. (The figure includes most of the world's economies. It excludes major oil-producing countries and countries that were communist during much of this period, because their experiences are explained by their special

figure 7-6



International Evidence on Investment Rates and Income per Person This scatterplot shows the experience of 84 countries, each represented by a single point. The horizontal axis shows the country's rate of investment, and the vertical axis shows the country's income per person. High investment is associated with high income per person, as the Solow model predicts.

Source: Robert Summers and Alan Heston, Supplement (Mark 5.6) to "The Penn World Table (Mark 5): An Expanded Set of International Comparisons 1950–1988," *Quarterly Journal of Economics* (May 1991): 327–368.

circumstances.) The data show a positive relationship between the fraction of output devoted to investment and the level of income per person. That is, countries with high rates of investment, such as the United States and Japan, usually have high incomes, whereas countries with low rates of investment, such as Uganda and Chad, have low incomes. Thus, the data are consistent with the Solow model's prediction that the investment rate is a key determinant of whether a country is rich or poor.

The strong correlation shown in this figure is an important fact, but it raises as many questions as it resolves. One might naturally ask, why do rates of saving and investment vary so much from country to country? There are many potential answers, such as tax policy, retirement patterns, the development of financial markets, and cultural differences. In addition, political stability may play a role: not surprisingly, rates of saving and investment tend to be low in countries with frequent wars, revolutions, and coups. Saving and investment also tend to be low in countries with poor political institutions, as measured by estimates of official corruption. A final interpretation of the evidence in Figure 7-6 is reverse causation:

perhaps high levels of income somehow foster high rates of saving and investment. Unfortunately, there is no consensus among economists about which of the many possible explanations is most important.

The association between investment rates and income per person is strong, and it is an important clue as to why some countries are rich and others poor, but it is not the whole story. The correlation between these two variables is far from perfect. Mexico and Zimbabwe, for instance, have had similar investment rates, but income per person is more than three times higher in Mexico. There must be other determinants of living standards beyond saving and investment. We therefore return to the international differences later in the chapter to see what other variables enter the picture.

7-2 The Golden Rule Level of Capital

So far, we have used the Solow model to examine how an economy's rate of saving and investment determines its steady-state levels of capital and income. This analysis might lead you to think that higher saving is always a good thing, for it always leads to greater income. Yet suppose a nation had a saving rate of 100 percent. That would lead to the largest possible capital stock and the largest possible income. But if all of this income is saved and none is ever consumed, what good is it?

This section uses the Solow model to discuss what amount of capital accumulation is optimal from the standpoint of economic well-being. In the next chapter, we discuss how government policies influence a nation's saving rate. But first, in this section, we present the theory behind these policy decisions.

Comparing Steady States

To keep our analysis simple, let's assume that a policymaker can set the economy's saving rate at any level. By setting the saving rate, the policymaker determines the economy's steady state. What steady state should the policymaker choose?

When choosing a steady state, the policymaker's goal is to maximize the well-being of the individuals who make up the society. Individuals themselves do not care about the amount of capital in the economy, or even the amount of output. They care about the amount of goods and services they can consume. Thus, a benevolent policymaker would want to choose the steady state with the highest level of consumption. The steady-state value of k that maximizes consumption is called the **Golden Rule level of capital** and is denoted k_{gold}^* .³

How can we tell whether an economy is at the Golden Rule level? To answer this question, we must first determine steady-state consumption per worker. Then we can see which steady state provides the most consumption.

³ Edmund Phelps, "The Golden Rule of Accumulation: A Fable for Growthmen," *American Economic Review* 51 (September 1961): 638–643.

To find steady-state consumption per worker, we begin with the national income accounts identity

$$y = c + i$$

and rearrange it as

$$c = y - i.$$


Consumption is simply output minus investment. Because we want to find steady-state consumption, we substitute steady-state values for output and investment. Steady-state output per worker is $f(k^*)$, where k^* is the steady-state capital stock per worker. Furthermore, because the capital stock is not changing in the steady state, investment is equal to depreciation δk^* . Substituting $f(k^*)$ for y and δk^* for i , we can write steady-state consumption per worker as

$$c^* = f(k^*) - \delta k^*.$$

According to this equation, steady-state consumption is what's left of steady-state output after paying for steady-state depreciation. This equation shows that an increase in steady-state capital has two opposing effects on steady-state consumption. On the one hand, more capital means more output. On the other hand, more capital also means that more output must be used to replace capital that is wearing out.

Figure 7-7 graphs steady-state output and steady-state depreciation as a function of the steady-state capital stock. Steady-state consumption is the gap between

figure 7-7



Steady-State Consumption

The economy's output is used for consumption or investment. In the steady state, investment equals depreciation. Therefore, steady-state consumption is the difference between output $f(k^*)$ and depreciation δk^* . Steady-state consumption is maximized at the Golden Rule steady state. The Golden Rule capital stock is denoted k^*_{gold} , and the Golden Rule consumption is denoted c^*_{gold} .

output and depreciation. This figure shows that there is one level of the capital stock—the Golden Rule level k_{gold}^* —that maximizes consumption.

When comparing steady states, we must keep in mind that higher levels of capital affect both output and depreciation. If the capital stock is below the Golden Rule level, an increase in the capital stock raises output more than depreciation, so that consumption rises. In this case, the production function is steeper than the δk^* line, so the gap between these two curves—which equals consumption—grows as k^* rises. By contrast, if the capital stock is above the Golden Rule level, an increase in the capital stock reduces consumption, since the increase in output is smaller than the increase in depreciation. In this case, the production function is flatter than the δk^* line, so the gap between the curves—consumption—shrinks as k^* rises. At the Golden Rule level of capital, the production function and the δk^* line have the same slope, and consumption is at its greatest level.

We can now derive a simple condition that characterizes the Golden Rule level of capital. Recall that the slope of the production function is the marginal product of capital MPK . The slope of the δk^* line is δ . Because these two slopes are equal at k_{gold}^* , the Golden Rule is described by the equation

$$MPK = \delta.$$

At the Golden Rule level of capital, the marginal product of capital equals the depreciation rate.

To make the point somewhat differently, suppose that the economy starts at some steady-state capital stock k^* and that the policymaker is considering increasing the capital stock to $k^* + 1$. The amount of extra output from this increase in capital would be $f(k^* + 1) - f(k^*)$, which is the marginal product of capital MPK . The amount of extra depreciation from having 1 more unit of capital is the depreciation rate δ . Thus, the net effect of this extra unit of capital on consumption is $MPK - \delta$. If $MPK - \delta > 0$, then increases in capital increase consumption, so k^* must be below the Golden Rule level. If $MPK - \delta < 0$, then increases in capital decrease consumption, so k^* must be above the Golden Rule level. Therefore, the following condition describes the Golden Rule:

$$MPK - \delta = 0.$$


At the Golden Rule level of capital, the marginal product of capital net of depreciation ($MPK - \delta$) equals zero. As we will see, a policymaker can use this condition to find the Golden Rule capital stock for an economy.⁴

Keep in mind that the economy does not automatically gravitate toward the Golden Rule steady state. If we want any particular steady-state capital stock, such as the Golden Rule, we need a particular saving rate to support it. Figure 7–8

⁴ *Mathematical note:* Another way to derive the condition for the Golden Rule uses a bit of calculus. Recall that $c^* = f(k^*) - \delta k^*$. To find the k^* that maximizes c^* , differentiate to find $dc^*/dk^* = f'(k^*) - \delta$ and set this derivative equal to zero. Noting that $f'(k^*)$ is the marginal product of capital, we obtain the Golden Rule condition in the text.

figure 7-8

Steady-state output,
depreciation, and
investment per worker



1. To reach the
Golden Rule
steady state...

2. ...the economy
needs the right
saving rate.

The Saving Rate and the Golden Rule There is only one saving rate that produces the Golden Rule level of capital k^*_{gold} . Any change in the saving rate would shift the $sf(k)$ curve and would move the economy to a steady state with a lower level of consumption.

shows the steady state if the saving rate is set to produce the Golden Rule level of capital. If the saving rate is higher than the one used in this figure, the steady-state capital stock will be too high. If the saving rate is lower, the steady-state capital stock will be too low. In either case, steady-state consumption will be lower than it is at the Golden Rule steady state.

Finding the Golden Rule Steady State: A Numerical Example

Consider the decision of a policymaker choosing a steady state in the following economy. The production function is the same as in our earlier example:

$$y = \sqrt{k}.$$

Output per worker is the square root of capital per worker. Depreciation δ is again 10 percent of capital. This time, the policymaker chooses the saving rate s and thus the economy's steady state.

To see the outcomes available to the policymaker, recall that the following equation holds in the steady state:

$$\frac{k^*}{f(k^*)} = \frac{s}{\delta}.$$

In this economy, this equation becomes

$$\frac{k^*}{\sqrt{k^*}} = \frac{s}{0.1}.$$

Squaring both sides of this equation yields a solution for the steady-state capital stock. We find

$$k^* = 100s^2.$$

Using this result, we can compute the steady-state capital stock for any saving rate.

Table 7-3 presents calculations showing the steady states that result from various saving rates in this economy. We see that higher saving leads to a higher capital stock, which in turn leads to higher output and higher depreciation. Steady-state consumption, the difference between output and depreciation, first rises with higher saving rates and then declines. Consumption is highest when the saving rate is 0.5. Hence, a saving rate of 0.5 produces the Golden Rule steady state.

table 7-3

Finding the Golden Rule Steady State: A Numerical Example

Assumptions: $y = \sqrt{k}$; $\delta = 0.1$

s	k^*	y^*	δk^*	c^*	MPK	$MPK - \delta$
0.0	0.0	0.0	0.0	0.0	∞	∞
0.1	1.0	1.0	0.1	0.9	0.500	0.400
0.2	4.0	2.0	0.4	1.6	0.250	0.150
0.3	9.0	3.0	0.9	2.1	0.167	0.067
0.4	16.0	4.0	1.6	2.4	0.125	0.025
0.5	25.0	5.0	2.5	2.5	0.100	0.000
0.6	36.0	6.0	3.6	2.4	0.083	-0.017
0.7	49.0	7.0	4.9	2.1	0.071	-0.029
0.8	64.0	8.0	6.4	1.6	0.062	-0.038
0.9	81.0	9.0	8.1	0.9	0.056	-0.044
1.0	100.0	10.0	10.0	0.0	0.050	-0.050

Recall that another way to identify the Golden Rule steady state is to find the capital stock at which the net marginal product of capital ($MPK - \delta$) equals zero. For this production function, the marginal product is⁵

$$MPK = \frac{1}{2\sqrt{k}}.$$

⁵ *Mathematical note:* To derive this formula, note that the marginal product of capital is the derivative of the production function with respect to k .

Using this formula, the last two columns of Table 7-3 present the values of MPK and $MPK - \delta$ in the different steady states. Note that the net marginal product of capital is exactly zero when the saving rate is at its Golden Rule value of 0.5. Because of diminishing marginal product, the net marginal product of capital is greater than zero whenever the economy saves less than this amount, and it is less than zero whenever the economy saves more.

This numerical example confirms that the two ways of finding the Golden Rule steady state—looking at steady-state consumption or looking at the marginal product of capital—give the same answer. If we want to know whether an actual economy is currently at, above, or below its Golden Rule capital stock, the second method is usually more convenient, because estimates of the marginal product of capital are easy to come by. By contrast, evaluating an economy with the first method requires estimates of steady-state consumption at many different saving rates; such information is hard to obtain. Thus, when we apply this kind of analysis to the U.S. economy in the next chapter, we will find it useful to examine estimates of the marginal product of capital.

The Transition to the Golden Rule Steady State


Let's now make our policymaker's problem more realistic. So far, we have been assuming that the policymaker can simply choose the economy's steady state and jump there immediately. In this case, the policymaker would choose the steady state with highest consumption—the Golden Rule steady state. But now suppose that the economy has reached a steady state other than the Golden Rule. What happens to consumption, investment, and capital when the economy makes the transition between steady states? Might the impact of the transition deter the policymaker from trying to achieve the Golden Rule?

We must consider two cases: the economy might begin with more capital than in the Golden Rule steady state, or with less. It turns out that the two cases offer very different problems for policymakers. (As we will see in the next chapter, the second case—too little capital—describes most actual economies, including that of the United States.)

Starting With Too Much Capital We first consider the case in which the economy begins at a steady state with more capital than it would have in the Golden Rule steady state. In this case, the policymaker should pursue policies aimed at reducing the rate of saving in order to reduce the capital stock. Suppose that these policies succeed and that at some point—call it time t_0 —the saving rate falls to the level that will eventually lead to the Golden Rule steady state.

Figure 7-9 shows what happens to output, consumption, and investment when the saving rate falls. The reduction in the saving rate causes an immediate increase in consumption and a decrease in investment. Because investment and depreciation were equal in the initial steady state, investment will now be less than depreciation, which means the economy is no longer in a steady state. Gradually, the capital stock falls, leading to reductions in output, consumption, and investment. These variables continue to fall until the economy reaches the new steady state. Because we are assuming that the new steady state is the Golden

figure 7-9



Reducing Saving When Starting With More Capital Than in the Golden Rule Steady State

This figure shows what happens over time to output, consumption, and investment when the economy begins with more capital than the Golden Rule level and the saving rate is reduced. The reduction in the saving rate (at time t_0) causes an immediate increase in consumption and an equal decrease in investment. Over time, as the capital stock falls, output, consumption, and investment fall together. Because the economy began with too much capital, the new steady state has a higher level of consumption than the initial steady state.

Rule steady state, consumption must be higher than it was before the change in the saving rate, even though output and investment are lower.


Note that, compared to the old steady state, consumption is higher not only in the new steady state but also along the entire path to it. When the capital stock exceeds the Golden Rule level, reducing saving is clearly a good policy, for it increases consumption at every point in time.

Starting With Too Little Capital When the economy begins with less capital than in the Golden Rule steady state, the policymaker must raise the saving rate to reach the Golden Rule. Figure 7-10 shows what happens. The increase in the saving rate at time t_0 causes an immediate fall in consumption and a rise in investment. Over time, higher investment causes the capital stock to rise. As capital accumulates, output, consumption, and investment gradually increase, eventually approaching the new steady-state levels. Because the initial steady state was below the Golden Rule, the increase in saving eventually leads to a higher level of consumption than that which prevailed initially.

Does the increase in saving that leads to the Golden Rule steady state raise economic welfare? Eventually it does, because the steady-state level of consumption is higher. But achieving that new steady state requires an initial period of reduced consumption. Note the contrast to the case in which the economy begins above the Golden Rule. *When the economy begins above the Golden Rule, reaching the Golden Rule produces higher consumption at all points in time. When the economy begins below the Golden Rule, reaching the Golden Rule requires initially reducing consumption to increase consumption in the future.*

When deciding whether to try to reach the Golden Rule steady state, policymakers have to take into account that current consumers and future consumers are not always the same people. Reaching the Golden Rule achieves the highest

figure 7-10



Increasing Saving When Starting With Less Capital Than in the Golden Rule Steady State

This figure shows what happens over time to output, consumption, and investment when the economy begins with less capital than the Golden Rule, and the saving rate is increased. The increase in the saving rate (at time t_0) causes an immediate drop in consumption and an equal jump in investment. Over time, as the capital stock grows, output, consumption, and investment increase together. Because the economy began with less capital than the Golden Rule, the new steady state has a higher level of consumption than the initial steady state.

steady-state level of consumption and thus benefits future generations. But when the economy is initially below the Golden Rule, reaching the Golden Rule requires raising investment and thus lowering the consumption of current generations. Thus, when choosing whether to increase capital accumulation, the policymaker faces a tradeoff among the welfare of different generations. A policymaker who cares more about current generations than about future generations may decide not to pursue policies to reach the Golden Rule steady state. By contrast, a policymaker who cares about all generations equally will choose to reach the Golden Rule. Even though current generations will consume less, an infinite number of future generations will benefit by moving to the Golden Rule.

Thus, optimal capital accumulation depends crucially on how we weigh the interests of current and future generations. The biblical Golden Rule tells us, “do unto others as you would have them do unto you.” If we heed this advice, we give all generations equal weight. In this case, it is optimal to reach the Golden Rule level of capital—which is why it is called the “Golden Rule.”

7-3 Population Growth

The basic Solow model shows that capital accumulation, by itself, cannot explain sustained economic growth: high rates of saving lead to high growth temporarily, but the economy eventually approaches a steady state in which capital and output are constant. To explain the sustained economic growth that we observe in most parts of the world, we must expand the Solow model to incorporate the other two sources of economic growth—population growth and technological progress. In this section we add population growth to the model.

Instead of assuming that the population is fixed, as we did in Sections 7-1 and 7-2, we now suppose that the population and the labor force grow at a constant rate n . For example, the U.S. population grows about 1 percent per year, so $n = 0.01$. This means that if 150 million people are working one year, then 151.5 million (1.01×150) are working the next year, and 153.015 million (1.01×151.5) the year after that, and so on.

The Steady State With Population Growth

How does population growth affect the steady state? To answer this question, we must discuss how population growth, along with investment and depreciation, influences the accumulation of capital per worker. As we noted before, investment raises the capital stock, and depreciation reduces it. But now there is a third force acting to change the amount of capital per worker: the growth in the number of workers causes capital per worker to fall.

We continue to let lowercase letters stand for quantities per worker. Thus, $k = K/L$ is capital per worker, and $y = Y/L$ is output per worker. Keep in mind, however, that the number of workers is growing over time.

The change in the capital stock per worker is

$$\Delta k = i - (\delta + n)k.$$

This equation shows how investment, depreciation, and population growth influence the per-worker capital stock. Investment increases k , whereas depreciation and population growth decrease k . We saw this equation earlier in this chapter for the special case of a constant population ($n = 0$).


We can think of the term $(\delta + n)k$ as defining *break-even investment*—the amount of investment necessary to keep the capital stock per worker constant. Break-even investment includes the depreciation of existing capital, which equals δk . It also includes the amount of investment necessary to provide new workers with capital. The amount of investment necessary for this purpose is nk , because there are n new workers for each existing worker, and because k is the amount of capital for each worker. The equation shows that population growth reduces the accumulation of capital per worker much the way depreciation does. Depreciation reduces k by wearing out the capital stock, whereas population growth reduces k by spreading the capital stock more thinly among a larger population of workers.⁶

Our analysis with population growth now proceeds much as it did previously. First, we substitute $sf(k)$ for i . The equation can then be written as

$$\Delta k = sf(k) - (\delta + n)k.$$

⁶ *Mathematical note:* Formally deriving the equation for the change in k requires a bit of calculus. Note that the change in k per unit of time is $dk/dt = d(K/L)/dt$. After applying the chain rule, we can write this as $dk/dt = (1/L)(dK/dt) - (K/L^2)(dL/dt)$. Now use the following facts to substitute in this equation: $dK/dt = I - \delta K$ and $(dL/dt)/L = n$. After a bit of manipulation, this produces the equation in the text.

figure 7-11



Population Growth in the Solow Model Like depreciation, population growth is one reason why the capital stock per worker shrinks. If n is the rate of population growth and δ is the rate of depreciation, then $(\delta + n)k$ is *break-even investment*—the amount of investment necessary to keep constant the capital stock per worker k . For the economy to be in a steady state, investment $sf(k)$ must offset the effects of depreciation and population growth $(\delta + n)k$. This is represented by the crossing of the two curves.

To see what determines the steady-state level of capital per worker, we use Figure 7-11, which extends the analysis of Figure 7-4 to include the effects of population growth. An economy is in a steady state if capital per worker k is unchanging. As before, we designate the steady-state value of k as k^* . If k is less than k^* , investment is greater than break-even investment, so k rises. If k is greater than k^* , investment is less than break-even investment, so k falls.


In the steady state, the positive effect of investment on the capital stock per worker exactly balances the negative effects of depreciation and population growth. That is, at k^* , $\Delta k = 0$ and $i^* = \delta k^* + nk^*$. Once the economy is in the steady state, investment has two purposes. Some of it (δk^*) replaces the depreciated capital, and the rest (nk^*) provides the new workers with the steady-state amount of capital.

The Effects of Population Growth

Population growth alters the basic Solow model in three ways. First, it brings us closer to explaining sustained economic growth. In the steady state with population growth, capital per worker and output per worker are constant. Because the number of workers is growing at rate n , however, *total* capital and *total* output must also be growing at rate n . Hence, although population growth cannot explain sustained growth in the standard of living (because output per worker is constant in the steady state), it can help explain sustained growth in total output.

Second, population growth gives us another explanation for why some countries are rich and others are poor. Consider the effects of an increase in population growth. Figure 7-12 shows that an increase in the rate of population growth from n_1 to n_2 reduces the steady-state level of capital per worker from k_1^* to k_2^* .

figure 7-12

**The Impact of Population Growth**

An increase in the rate of population growth from n_1 to n_2 shifts the line representing population growth and depreciation upward. The new steady state k_2^* has a lower level of capital per worker than the initial steady state k_1^* . Thus, the Solow model predicts that economies with higher rates of population growth will have lower levels of capital per worker and therefore lower incomes.

Because k^* is lower, and because $y^* = f(k^*)$, the level of output per worker y^* is also lower. Thus, the Solow model predicts that countries with higher population growth will have lower levels of GDP per person.

Finally, population growth affects our criterion for determining the Golden Rule (consumption-maximizing) level of capital. To see how this criterion changes, note that consumption per worker is

$$c = y - i.$$

Because steady-state output is $f(k^*)$ and steady-state investment is $(\delta + n)k^*$, we can express steady-state consumption as

$$c^* = f(k^*) - (\delta + n)k^*.$$

Using an argument largely the same as before, we conclude that the level of k^* that maximizes consumption is the one at which

$$MPK = \delta + n,$$

or equivalently,

$$MPK - \delta = n.$$

In the Golden Rule steady state, the marginal product of capital net of depreciation equals the rate of population growth.


CASE STUDY

Population Growth Around the World

Let's return now to the question of why standards of living vary so much around the world. The analysis we have just completed suggests that population growth may be one of the answers. According to the Solow model, a nation with a high rate of population growth will have a low steady-state capital stock per worker and thus also a low level of income per worker. In other words, high population growth tends to impoverish a country because it is hard to maintain a high level of capital per worker when the number of workers is growing quickly. To see whether the evidence supports this conclusion, we again look at cross-country data.

Figure 7-13 is a scatterplot of data for the same 84 countries examined in the previous case study (and in Figure 7-6). The figure shows that countries with high rates of population growth tend to have low levels of income per person. The international evidence is consistent with our model's prediction that the rate of population growth is one determinant of a country's standard of living.

figure 7-13



International Evidence on Population Growth and Income per Person This figure is a scatterplot of data from 84 countries. It shows that countries with high rates of population growth tend to have low levels of income per person, as the Solow model predicts.

Source: Robert Summers and Alan Heston, Supplement (Mark 5.6) to "The Penn World Table (Mark 5): An Expanded Set of International Comparisons 1950–1988," *Quarterly Journal of Economics* (May 1991): 327–368.

This conclusion is not lost on policymakers. Those trying to pull the world's poorest nations out of poverty, such as the advisers sent to developing nations by the World Bank, often advocate reducing fertility by increasing education about birth-control methods and expanding women's job opportunities. Toward the same end, China has followed the totalitarian policy of allowing only one child per couple. These policies to reduce population growth should, if the Solow model is right, raise income per person in the long run.

In interpreting the cross-country data, however, it is important to keep in mind that correlation does not imply causation. The data show that low population growth is typically associated with high levels of income per person, and the Solow model offers one possible explanation for this fact, but other explanations are also possible. It is conceivable that high income encourages low population growth, perhaps because birth-control techniques are more readily available in richer countries. The international data can help us evaluate a theory of growth, such as the Solow model, because they show us whether the theory's predictions are borne out in the world. But often more than one theory can explain the same facts.

7-4 Conclusion

This chapter has started the process of building the Solow growth model. The model as developed so far shows how saving and population growth determine the economy's steady-state capital stock and its steady-state level of income per person. As we have seen, it sheds light on many features of actual growth experiences—why Germany and Japan grew so rapidly after being devastated by World War II, why countries that save and invest a high fraction of their output are richer than countries that save and invest a smaller fraction, and why countries with high rates of population growth are poorer than countries with low rates of population growth.

What the model cannot do, however, is explain the persistent growth in living standards we observe in most countries. In the model we now have, when the economy reaches its steady state, output per worker stops growing. To explain persistent growth, we need to introduce technological progress into the model. That is our first job in the next chapter.

Summary

1. The Solow growth model shows that in the long run, an economy's rate of saving determines the size of its capital stock and thus its level of production. The higher the rate of saving, the higher the stock of capital and the higher the level of output.
2. In the Solow model, an increase in the rate of saving causes a period of rapid growth, but eventually that growth slows as the new steady state is reached.

Thus, although a high saving rate yields a high steady-state level of output, saving by itself cannot generate persistent economic growth.

3. The level of capital that maximizes steady-state consumption is called the Golden Rule level. If an economy has more capital than in the Golden Rule steady state, then reducing saving will increase consumption at all points in time. By contrast, if the economy has less capital in the Golden Rule steady state, then reaching the Golden Rule requires increased investment and thus lower consumption for current generations.
4. The Solow model shows that an economy's rate of population growth is another long-run determinant of the standard of living. The higher the rate of population growth, the lower the level of output per worker.

KEY CONCEPTS

Solow growth model

Steady state

Golden Rule level of capital

QUESTIONS FOR REVIEW

1. In the Solow model, how does the saving rate affect the steady-state level of income? How does it affect the steady-state rate of growth?
2. Why might an economic policymaker choose the Golden Rule level of capital?
3. Might a policymaker choose a steady state with more capital than in the Golden Rule steady state? With less capital than in the Golden Rule steady state? Explain your answers.
4. In the Solow model, how does the rate of population growth affect the steady-state level of income? How does it affect the steady-state rate of growth?

PROBLEMS AND APPLICATIONS

1. Country A and country B both have the production function

$$Y = F(K, L) = K^{1/2}L^{1/2}.$$

- a. Does this production function have constant returns to scale? Explain.
- b. What is the per-worker production function, $y = f(k)$?
- c. Assume that neither country experiences population growth or technological progress and that 5 percent of capital depreciates each year. Assume further that country A saves 10 percent of output each year and country B saves 20 percent of output each year. Using your answer from part (b) and the steady-state condition that investment equals depreciation, find the

steady-state level of capital per worker for each country. Then find the steady-state levels of income per worker and consumption per worker.

- d. Suppose that both countries start off with a capital stock per worker of 2. What are the levels of income per worker and consumption per worker? Remembering that the change in the capital stock is investment less depreciation, use a calculator to show how the capital stock per worker will evolve over time in both countries. For each year, calculate income per worker and consumption per worker. How many years will it be before the consumption in country B is higher than the consumption in country A?
2. In the discussion of German and Japanese postwar growth, the text describes what happens

when part of the capital stock is destroyed in a war. By contrast, suppose that a war does not directly affect the capital stock, but that casualties reduce the labor force.

- a. What is the immediate impact on total output and on output per person?
 - b. Assuming that the saving rate is unchanged and that the economy was in a steady state before the war, what happens subsequently to output per worker in the postwar economy? Is the growth rate of output per worker after the war smaller or greater than normal?
3. Consider an economy described by the production function $Y = F(K, L) = K^{0.3}L^{0.7}$.
- a. What is the per-worker production function?
 - b. Assuming no population growth or technological progress, find the steady-state capital stock per worker, output per worker, and consumption per worker as functions of the saving rate and the depreciation rate.
 - c. Assume that the depreciation rate is 10 percent per year. Make a table showing steady-state capital per worker, output per worker, and consumption per worker for saving rates of 0 percent, 10 percent, 20 percent, 30 percent, and so on. (You will need a calculator with an exponent key for this.) What saving rate maximizes output per worker? What saving rate maximizes consumption per worker?
 - d. (Harder) Use calculus to find the marginal product of capital. Add to your table the marginal product of capital net of depreciation for each of the saving rates. What does your table show?
4. The 1983 *Economic Report of the President* contained the following statement: "Devoting a larger share of national output to investment would help restore rapid productivity growth and rising living standards." Do you agree with this claim? Explain.
5. One view of the consumption function is that workers have high propensities to consume and capitalists have low propensities to consume. To explore the implications of this view, suppose that an economy consumes all wage income and saves all capital income. Show that if the factors of production earn their marginal product, this economy reaches the Golden Rule level of capital. (*Hint:* Begin with the identity that saving equals investment. Then use the steady-state condition that investment is just enough to keep up with depreciation and population growth, and the fact that saving equals capital income in this economy.)
6. Many demographers predict that the United States will have zero population growth in the twenty-first century, in contrast to average population growth of about 1 percent per year in the twentieth century. Use the Solow model to forecast the effect of this slowdown in population growth on the growth of total output and the growth of output per person. Consider the effects both in the steady state and in the transition between steady states.
7. In the Solow model, population growth leads to steady-state growth in total output, but not in output per worker. Do you think this would still be true if the production function exhibited increasing or decreasing returns to scale? Explain. (For the definitions of increasing and decreasing returns to scale, see Chapter 3, "Problems and Applications," Problem 2.)
8. Consider how unemployment would affect the Solow growth model. Suppose that output is produced according to the production function $Y = K^\alpha[(1 - u)L]^{1-\alpha}$, where K is capital, L is the labor force, and u is the natural rate of unemployment. The national saving rate is s , the labor force grows at rate n , and capital depreciates at rate δ .
- a. Express output per worker ($y = Y/L$) as a function of capital per worker ($k = K/L$) and the natural rate of unemployment. Describe the steady state of this economy.
 - b. Suppose that some change in government policy reduces the natural rate of unemployment. Describe how this change affects output both immediately and over time. Is the steady-state effect on output larger or smaller than the immediate effect? Explain.
9. Choose two countries that interest you—one rich and one poor. What is the income per person in each country? Find some data on country characteristics that might help explain the difference in income: investment rates, population growth rates, educational attainment, and so on. (*Hint:* The Web site of the World Bank, www.worldbank.org, is one place to find such data.) How might you figure out which of these factors is most responsible for the observed income difference?

Economic Growth II

Is there some action a government of India could take that would lead the Indian economy to grow like Indonesia's or Egypt's? If so, what, exactly? If not, what is it about the "nature of India" that makes it so? The consequences for human welfare involved in questions like these are simply staggering: Once one starts to think about them, it is hard to think about anything else.

— Robert E. Lucas, Jr.

This chapter continues our analysis of the forces governing long-run economic growth. With the basic version of the Solow growth model as our starting point, we take on four new tasks.

Our first task is to make the Solow model more general and more realistic. In Chapter 3 we saw that capital, labor, and technology are the key determinants of a nation's production of goods and services. In Chapter 7 we developed the Solow model to show how changes in capital (saving and investment) and changes in the labor force (population growth) affect the economy's output. We are now ready to add the third source of growth—changes in technology—into the mix.

Our second task is to examine how a nation's public policies can influence the level and growth of its standard of living. In particular, we address four questions: Should our society save more or save less? How can policy influence the rate of saving? Are there some types of investment that policy should especially encourage? How can policy increase the rate of technological progress? The Solow growth model provides the theoretical framework within which we consider each of these policy issues.

Our third task is to move from theory to empirics. That is, we consider how well the Solow model fits the facts. During the 1990s, a large literature examined the predictions of the Solow model and other models of economic growth. It turns out that the glass is both half full and half empty. The Solow model can shed much light on international growth experiences, but it is far from the last word on the subject.

Our fourth and final task is to consider what the Solow model leaves out. As we have discussed previously, models help us understand the world by simplifying it. After completing an analysis of a model, therefore, it is important to consider

whether we have oversimplified matters. In the last section, we examine a new set of theories, called *endogenous growth theories*, that hope to explain the technological progress that the Solow model takes as exogenous.

8-1 Technological Progress in the Solow Model

So far, our presentation of the Solow model has assumed an unchanging relationship between the inputs of capital and labor and the output of goods and services. Yet the model can be modified to include exogenous technological progress, which over time expands society's ability to produce.

The Efficiency of Labor

To incorporate technological progress, we must return to the production function that relates total capital K and total labor L to total output Y . Thus far, the production function has been

$$Y = F(K, L).$$

We now write the production function as

$$Y = F(K, L \times E),$$

where E is a new (and somewhat abstract) variable called the **efficiency of labor**. The efficiency of labor is meant to reflect society's knowledge about production methods: as the available technology improves, the efficiency of labor rises. For instance, the efficiency of labor rose when assembly-line production transformed manufacturing in the early twentieth century, and it rose again when computerization was introduced in the late twentieth century. The efficiency of labor also rises when there are improvements in the health, education, or skills of the labor force.

The term $L \times E$ measures the number of *effective workers*. It takes into account the number of workers L and the efficiency of each worker E . This new production function states that total output Y depends on the number of units of capital K and on the number of effective workers $L \times E$. Increases in the efficiency of labor E are, in effect, like increases in the labor force L .

The simplest assumption about technological progress is that it causes the efficiency of labor E to grow at some constant rate g . For example, if $g = 0.02$, then each unit of labor becomes 2 percent more efficient each year: output increases as if the labor force had increased by an additional 2 percent. This form of technological progress is called *labor augmenting*, and g is called the rate of **labor-augmenting technological progress**. Because the labor force L is growing at rate n , and the efficiency of each unit of labor E is growing at rate g , the number of effective workers $L \times E$ is growing at rate $n + g$.

The Steady State With Technological Progress

Expressing technological progress as labor augmenting makes it analogous to population growth. In Chapter 7 we analyzed the economy in terms of quantities per worker and allowed the number of workers to rise over time. Now we analyze the economy in terms of quantities per effective worker and allow the number of effective workers to rise.

To do this, we need to reconsider our notation. We now let $k = K/(L \times E)$ stand for capital per effective worker and $y = Y/(L \times E)$ stand for output per effective worker. With these definitions, we can again write $y = f(k)$.

This notation is not really as new as it seems. If we hold the efficiency of labor E constant at the arbitrary value of 1, as we have done implicitly up to now, then these new definitions of k and y reduce to our old ones. When the efficiency of labor is growing, however, we must keep in mind that k and y now refer to quantities per effective worker (not per actual worker).


Our analysis of the economy proceeds just as it did when we examined population growth. The equation showing the evolution of k over time now changes to

$$\Delta k = sf(k) - (\delta + n + g)k.$$

As before, the change in the capital stock Δk equals investment $sf(k)$ minus break-even investment $(\delta + n + g)k$. Now, however, because $k = K/EL$, break-even investment includes three terms: to keep k constant, δk is needed to replace depreciating capital, nk is needed to provide capital for new workers, and gk is needed to provide capital for the new “effective workers” created by technological progress.

As shown in Figure 8-1, the inclusion of technological progress does not substantially alter our analysis of the steady state. There is one level of k , denoted

figure 8-1



Technological Progress and the Solow Growth Model Labor-augmenting technological progress at rate g affects the Solow growth model in much the same way as did population growth at rate n . Now that k is defined as the amount of capital per effective worker, increases in the number of effective workers because of technological progress tend to decrease k . In the steady state, investment $sf(k)$ exactly offsets the reductions in k attributable to depreciation, population growth, and technological progress.

k^* , at which capital per effective worker and output per effective worker are constant. As before, this steady state represents the long-run equilibrium of the economy.

The Effects of Technological Progress

Table 8-1 shows how four key variables behave in the steady state with technological progress. As we have just seen, capital per effective worker k is constant in the steady state. Because $y = f(k)$, output per effective worker is also constant. Remember, though, that the efficiency of each actual worker is growing at rate g . Hence, output per worker ($Y/L = y \times E$) also grows at rate g . Total output [$Y = y \times (E \times L)$] grows at rate $n + g$.

With the addition of technological progress, our model can finally explain the sustained increases in standards of living that we observe. That is, we have shown that technological progress can lead to sustained growth in output per worker. By contrast, a high rate of saving leads to a high rate of growth only until the steady state is reached. Once the economy is in steady state, the rate of growth of output per worker depends only on the rate of technological progress. *According to the Solow model, only technological progress can explain persistently rising living standards.*

The introduction of technological progress also modifies the criterion for the Golden Rule. The Golden Rule level of capital is now defined as the steady state that maximizes consumption per effective worker. Following the same arguments that we have used before, we can show that steady-state consumption per effective worker is

$$c^* = f(k^*) - (\delta + n + g)k^*.$$

Steady-state consumption is maximized if

$$MPK = \delta + n + g,$$

or

$$MPK - \delta = n + g.$$

That is, at the Golden Rule level of capital, the net marginal product of capital, $MPK - \delta$, equals the rate of growth of total output, $n + g$. Because actual

table 8-1

Steady-State Growth Rates in the Solow Model With Technological Progress

Variable	Symbol	Steady-State Growth Rate
Capital per effective worker	$k = K/(E \times L)$	0
Output per effective worker	$y = Y/(E \times L) = f(k)$	0
Output per worker	$Y/L = y \times E$	g
Total output	$Y = y \times (E \times L)$	$n + g$

economies experience both population growth and technological progress, we must use this criterion to evaluate whether they have more or less capital than at the Golden Rule steady state.

8-2 Policies to Promote Growth

Having used the Solow model to uncover the relationships among the different sources of economic growth, we can now use the theory to help guide our thinking about economic policy.

Evaluating the Rate of Saving

According to the Solow growth model, how much a nation saves and invests is a key determinant of its citizens' standard of living. So let's begin our policy discussion with a natural question: Is the rate of saving in the U.S. economy too low, too high, or about right?

As we have seen, the saving rate determines the steady-state levels of capital and output. One particular saving rate produces the Golden Rule steady state, which maximizes consumption per worker and thus economic well-being. The Golden Rule provides the benchmark against which we can compare the U.S. economy.

To decide whether the U.S. economy is at, above, or below the Golden Rule steady state, we need to compare the marginal product of capital net of depreciation ($MPK - \delta$) with the growth rate of total output ($n + g$). As we just established, at the Golden Rule steady state, $MPK - \delta = n + g$. If the economy is operating with less capital than in the Golden Rule steady state, then diminishing marginal product tells us that $MPK - \delta > n + g$. In this case, increasing the rate of saving will eventually lead to a steady state with higher consumption. However, if the economy is operating with too much capital, then $MPK - \delta < n + g$, and the rate of saving should be reduced.

To make this comparison for a real economy, such as the U.S. economy, we need an estimate of the growth rate ($n + g$) and an estimate of the net marginal product of capital ($MPK - \delta$). Real GDP in the United States grows an average of 3 percent per year, so $n + g = 0.03$. We can estimate the net marginal product of capital from the following three facts:

1. The capital stock is about 2.5 times one year's GDP.
2. Depreciation of capital is about 10 percent of GDP.
3. Capital income is about 30 percent of GDP.

Using the notation of our model (and the result from Chapter 3 that capital owners earn income of MPK for each unit of capital), we can write these facts as

1. $k = 2.5y$.
2. $\delta k = 0.1y$.
3. $MPK \times k = 0.3y$

We solve for the rate of depreciation δ by dividing equation 2 by equation 1:

$$\begin{aligned}\delta k/k &= (0.1\gamma)/(2.5\gamma) \\ \delta &= 0.04.\end{aligned}$$

And we solve for the marginal product of capital MPK by dividing equation 3 by equation 1:

$$\begin{aligned}(MPK \times k)/k &= (0.3\gamma)/(2.5\gamma) \\ MPK &= 0.12\end{aligned}$$

Thus, about 4 percent of the capital stock depreciates each year, and the marginal product of capital is about 12 percent per year. The net marginal product of capital, $MPK - \delta$, is about 8 percent per year.

We can now see that the return to capital ($MPK - \delta = 8$ percent per year) is well in excess of the economy's average growth rate ($n + g = 3$ percent per year). This fact, together with our previous analysis, indicates that the capital stock in the U.S. economy is well below the Golden Rule level. In other words, if the United States saved and invested a higher fraction of its income, it would grow more rapidly and eventually reach a steady state with higher consumption. This finding suggests that policymakers should want to increase the rate of saving and investment. In fact, for many years, increasing capital formation has been a high priority of economic policy.

Changing the Rate of Saving

The preceding calculations show that to move the U.S. economy toward the Golden Rule steady state, policymakers should increase national saving. But how can they do that? We saw in Chapter 3 that, as a matter of sheer accounting, higher national saving means higher public saving, higher private saving, or some combination of the two. Much of the debate over policies to increase growth centers on which of these options is likely to be most effective.

The most direct way in which the government affects national saving is through public saving—the difference between what the government receives in tax revenue and what it spends. When the government's spending exceeds its revenue, the government is said to run a *budget deficit*, which represents negative public saving. As we saw in Chapter 3, a budget deficit raises interest rates and crowds out investment; the resulting reduction in the capital stock is part of the burden of the national debt on future generations. Conversely, if the government spends less than it raises in revenue, it is said to run a *budget surplus*. It can then retire some of the national debt and stimulate investment.

The government also affects national saving by influencing private saving—the saving done by households and firms. In particular, how much people decide to save depends on the incentives they face, and these incentives are altered by a variety of public policies. Many economists argue that high tax rates on capital—including the corporate income tax, the federal income tax, the estate tax, and many state income and estate taxes—discourage private saving by reducing the

rate of return that savers earn. However, tax-exempt retirement accounts, such as IRAs, are designed to encourage private saving by giving preferential treatment to income saved in these accounts.

Many disagreements among economists over public policy are rooted in different views about how much private saving responds to incentives. For example, suppose that the government were to expand the amount that people could put into tax-exempt retirement accounts. Would people respond to the increased incentive to save by saving more? Or would people merely transfer saving done in other forms into these accounts—reducing tax revenue and thus public saving without any stimulus to private saving? Clearly, the desirability of the policy depends on the answers to these questions. Unfortunately, despite much research on this issue, no consensus has emerged.

CASE STUDY

Should the Social Security System Be Reformed?

Although many government policies are designed to encourage saving, such as the preferential tax treatment given to pension plans and other retirement accounts, one important policy is often thought to reduce saving: the Social Security system. Social Security is a transfer system designed to maintain individuals' income in their old age. These transfers to the elderly are financed with a payroll tax on the working-age population. This system is thought to reduce private saving because it reduces individuals' need to provide for their own retirement.

To counteract the reduction in national saving attributed to Social Security, many economists have proposed reforms of the Social Security system. The system is now largely *pay-as-you-go*: most of the current tax receipts are paid out to the current elderly population. One suggestion is that Social Security should be *fully funded*. Under this plan, the government would put aside in a trust fund the payments a generation makes when it is young and working; the government would then pay out the principal and accumulated interest to this same generation when it is older and retired. Under a fully funded Social Security system, an increase in public saving would offset the reduction in private saving.

A closely related proposal is *privatization*, which means turning this government program for the elderly into a system of mandatory private savings accounts, much like private pension plans. In principle, the issues of funding and privatization are distinct. A fully funded system could be either public (in which case the government holds the funds) or private (in which case private financial institutions hold the funds). In practice, however, the issues are often linked. Some economists have argued that a fully funded public system is problematic. They note that such a system would end up holding a large share of the nation's wealth, which would increase the role of the government in allocating capital. In addition, they fear that a large publicly controlled fund would tempt politicians to cut taxes or increase spending, which could deplete the fund and cause the system to revert to pay-as-you-go status. History gives some support to this fear: the initial architects of Social Security wanted the system to accumulate a much larger trust fund than ever materialized.

These issues rose to prominence in the late 1990s as policymakers became aware that the current Social Security system was not sustainable. That is, the amount of revenue being raised by the payroll tax appeared insufficient to pay all the benefits being promised. According to most projections, this problem was to become acute as the large baby-boom generation retired during the early decades of the twenty-first century. Various solutions were proposed. One possibility was to maintain the current system with some combination of smaller benefits and higher taxes. Other possibilities included movements toward a fully funded system, perhaps also including private accounts. This issue was prominent in the presidential campaign of 2000, with candidate George W. Bush advocating a reform including private accounts. As this book was going to press, it was still unclear whether this reform would come to pass.¹

Allocating the Economy's Investment

The Solow model makes the simplifying assumption that there is only one type of capital. In the world, of course, there are many types. Private businesses invest in traditional types of capital, such as bulldozers and steel plants, and newer types of capital, such as computers and robots. The government invests in various forms of public capital, called *infrastructure*, such as roads, bridges, and sewer systems.

In addition, there is *human capital*—the knowledge and skills that workers acquire through education, from early childhood programs such as Head Start to on-the-job training for adults in the labor force. Although the basic Solow model includes only physical capital and does not try to explain the efficiency of labor, in many ways human capital is analogous to physical capital. Like physical capital, human capital raises our ability to produce goods and services. Raising the level of human capital requires investment in the form of teachers, libraries, and student time. Recent research on economic growth has emphasized that human capital is at least as important as physical capital in explaining international differences in standards of living.²

Policymakers trying to stimulate economic growth must confront the issue of what kinds of capital the economy needs most. In other words, what kinds of capital yield the highest marginal products? To a large extent, policymakers can rely on the marketplace to allocate the pool of saving to alternative types of investment. Those industries with the highest marginal products of capital will naturally be most willing to borrow at market interest rates to finance new investment. Many economists advocate that the government should merely create a “level playing field” for different types of capital—for example, by ensuring that the tax system treats all forms of capital equally. The government can then rely on the market to allocate capital efficiently.

¹ To learn more about the debate over Social Security, see *Social Security Reform: Links to Saving, Investment, and Growth*, Steven A. Sass and Robert K. Triest, eds., Conference Series No. 41, Federal Reserve Bank of Boston, June 1997.

² N. Gregory Mankiw, David Romer, and David N. Weil, “A Contribution to the Empirics of Economic Growth,” *Quarterly Journal of Economics* (May 1992): 407–437.

Other economists have suggested that the government should actively encourage particular forms of capital. Suppose, for instance, that technological advance occurs as a by-product of certain economic activities. This would happen if new and improved production processes are devised during the process of building capital (a phenomenon called *learning by doing*) and if these ideas become part of society's pool of knowledge. Such a by-product is called a *technological externality* (or a *knowledge spillover*). In the presence of such externalities, the social returns to capital exceed the private returns, and the benefits of increased capital accumulation to society are greater than the Solow model suggests.³ Moreover, some types of capital accumulation may yield greater externalities than others. If, for example, installing robots yields greater technological externalities than building a new steel mill, then perhaps the government should use the tax laws to encourage investment in robots. The success of such an *industrial policy*, as it is sometimes called, requires that the government be able to measure the externalities of different economic activities so it can give the correct incentive to each activity.

Most economists are skeptical about industrial policies, for two reasons. First, measuring the externalities from different sectors is so difficult as to be virtually impossible. If policy is based on poor measurements, its effects might be close to random and, thus, worse than no policy at all. Second, the political process is far from perfect. Once the government gets in the business of rewarding specific industries with subsidies and tax breaks, the rewards are as likely to be based on political clout as on the magnitude of externalities.

One type of capital that necessarily involves the government is public capital. Local, state, and federal governments are always deciding whether to borrow to finance new roads, bridges, and transit systems. During his first presidential campaign, Bill Clinton argued that the United States had been investing too little in infrastructure. He claimed that a higher level of infrastructure investment would make the economy substantially more productive. Among economists, this claim had both defenders and critics. Yet all of them agree that measuring the marginal product of public capital is difficult. Private capital generates an easily measured rate of profit for the firm owning the capital, whereas the benefits of public capital are more diffuse.

Encouraging Technological Progress

The Solow model shows that sustained growth in income per worker must come from technological progress. The Solow model, however, takes technological progress as exogenous; it does not explain it. Unfortunately, the determinants of technological progress are not well understood.

Despite this limited understanding, many public policies are designed to stimulate technological progress. Most of these policies encourage the private sector to devote resources to technological innovation. For example, the patent system

³ Paul Romer, "Crazy Explanations for the Productivity Slowdown," *NBER Macroeconomics Annual* 2 (1987): 163–201.

gives a temporary monopoly to inventors of new products; the tax code offers tax breaks for firms engaging in research and development; and government agencies such as the National Science Foundation directly subsidize basic research in universities. In addition, as discussed above, proponents of industrial policy argue that the government should take a more active role in promoting specific industries that are key for rapid technological progress.

CASE STUDY

The Worldwide Slowdown in Economic Growth

Beginning in the early 1970s, world policymakers faced a perplexing problem—a global slowdown in economic growth. Table 8-2 presents data on the growth in real GDP per person for the seven major world economies. Growth in the United States fell from 2.2 percent to 1.5 percent, and other countries experienced similar or more severe declines. Accumulated over many years, even a small change in the rate of growth has a large effect on economic well-being. Real income in the United States today is about 20 percent lower than it would have been had growth remained at its previous level.

Why did this slowdown occur? Studies have shown that it was attributable to a fall in the rate at which the production function was improving over time. The appendix to this chapter explains how economists measure changes in the production function with a variable called *total factor productivity*, which is closely related to the efficiency of labor in the Solow model. There are, however, many hypotheses to explain this fall in productivity growth. Here are four of them.

Measurement Problems One possibility is that the productivity slowdown did not really occur and that it shows up in the data because the data are flawed. As you may recall from Chapter 2, one problem in measuring inflation is correcting for changes in the quality of goods and services. The same issue arises when measuring output and productivity. For instance, if technological advance leads to *more* computers being built, then the increase in output and productivity is easy to measure. But if technological advance leads to *faster* computers being built, then output and productivity have increased, but that increase is more subtle and harder to measure. Government statisticians try to correct for changes in quality, but despite their best efforts, the resulting data are far from perfect.

Unmeasured quality improvements mean that our standard of living is rising more rapidly than the official data indicate. This issue should make us suspicious of the data, but by itself it cannot explain the productivity slowdown. To explain a *slowdown* in growth, one must argue that the measurement problems got *worse*. There is some indication that this might be so. As history passes, fewer people work in industries with tangible and easily measured output, such as agriculture, and more work in industries with intangible and less easily measured output, such as medical services. Yet few economists believe that measurement problems were the full story.

Oil Prices When the productivity slowdown began around 1973, the obvious hypothesis to explain it was the large increase in oil prices caused by the actions of the OPEC oil cartel. The primary piece of evidence was the timing: productivity growth slowed at the same time that oil prices skyrocketed. Over time, however,

table 8-2

The Slowdown in Growth Around the World

Country	GROWTH IN OUTPUT PER PERSON (PERCENT PER YEAR)		
	1948–1972	1972–1995	1995–2000
Canada	2.9	1.8	2.7
France	4.3	1.6	2.2
West Germany	5.7	2.0	
Germany			1.7
Italy	4.9	2.3	4.7
Japan	8.2	2.6	1.1
United Kingdom	2.4	1.8	2.5
United States	2.2	1.5	2.9

Source: Angus Maddison, *Phases of Capitalist Development* (Oxford: Oxford University Press, 1982); and OECD National Accounts and International Financial Statistics.

Note: Data before 1995 for Germany refer to West Germany; after 1995, to the unified Germany.

this explanation has appeared less likely. One reason is that the accumulated shortfall in productivity seems too large to be explained by an increase in oil prices—oil is not that large a fraction of the typical firm's costs. In addition, if this explanation were right, productivity should have sped up when political turmoil in OPEC caused oil prices to plummet in 1986. Unfortunately, that did not happen.

Worker Quality Some economists suggest that the productivity slowdown might have been caused by changes in the labor force. In the early 1970s, the large baby-boom generation started leaving school and taking jobs. At the same time, changing social norms encouraged many women to leave full-time housework and enter the labor force. Both of these developments lowered the average level of experience among workers, which in turn lowered average productivity.

Other economists point to changes in worker quality as gauged by human capital. Although the educational attainment of the labor force continued to rise throughout this period, it was not increasing as rapidly as it did in the past. Moreover, declining performance on some standardized tests suggests that the quality of education was declining. If so, this could explain slowing productivity growth.

The Depletion of Ideas Still other economists suggest that the world started to run out of new ideas about how to produce in the early 1970s, pushing the economy into an age of slower technological progress. These economists often argue that the anomaly is not the period since 1970 but the preceding two decades. In the late 1940s, the economy had a large backlog of ideas that had not been fully implemented because of the Great Depression of the 1930s and World War II in the first half of 1940s. After the economy used up this backlog, the argument goes, a slowdown in productivity growth was likely. Indeed, although the growth rates in the 1970s, 1980s, and early 1990s were disappointing compared to those of the 1950s and 1960s, they were not lower than average growth rates from 1870 to 1950.

Which of these suspects is the culprit? All of them are plausible, but it is difficult to prove beyond a reasonable doubt that any one of them is guilty. The worldwide slowdown in economic growth that began in the mid-1970s remains a mystery.⁴

CASE STUDY

Information Technology and the New Economy

As any good doctor will tell you, sometimes a patient's illness goes away on its own, even if the doctor has failed to come up with a convincing diagnosis and remedy. This seems to be the outcome with the productivity slowdown discussed in the previous case study. Economists have not yet figured it out, but beginning in the middle of the 1990s, the problem disappeared. Economic growth took off, as shown in the third column of Table 8-2. In the United States, output per person accelerated from 1.5 to 2.9 percent per year. Commentators proclaimed we were living in a "new economy."

As with the slowdown in economic growth in the 1970s, the acceleration in the 1990s is hard to explain definitively. But part of the credit goes to advances in computer and information technology, including the Internet.

Observers of the computer industry often cite Moore's law, which states that the price of computing power falls by half every 18 months. This is not an inevitable law of nature but an empirical regularity describing the rapid technological progress this industry has enjoyed. In the 1980s and early 1990s, economists were surprised that the rapid progress in computing did not have a larger effect on the overall economy. Economist Robert Solow once quipped that "we can see the computer age everywhere but in the productivity statistics."

There are two reasons why the macroeconomic effects of the computer revolution might not have showed up until the mid-1990s. One is that the computer industry was previously only a small part of the economy. In 1990, computer hardware and software represented 0.9 percent of real GDP; by 1999, this share had risen to 4.2 percent. As computers made up a larger part of the economy, technological advance in that sector had a greater overall effect.

The second reason why the productivity benefits of computers may have been delayed is that it took time for firms to figure out how best to use the technology. Whenever firms change their production systems and train workers to use a technology, they disrupt the existing means of production. Measured productivity can fall for a while before the economy reaps the benefits. Indeed, some economists even suggest that the spread of computers can help explain the productivity slowdown that began in the 1970s.

Economic history provides some support for the idea that new technologies influence growth with a long lag. The electric light bulb was invented in 1879. But it took several decades before electricity had a big economic impact. For businesses to reap large productivity gains, they had to do more than simply re-

⁴ For various views on the growth slowdown, see "Symposium: The Slowdown in Productivity Growth," *The Journal of Economic Perspectives* 2 (Fall 1988): 3–98.

place steam engines with electric motors; they had to rethink the entire organization of factories. Similarly, replacing the typewriters on desks with computers and word processing programs, as was common in the 1980s, may have had small productivity effects. Only later, when the Internet and other advanced applications were invented, did the computers yield large economic gains.

Eventually, advances in technology should show up in economic growth, as was the case in the second half of the 1990s. This extra growth occurs through three channels. First, because the computer industry is part of the economy, productivity growth in that industry directly affects overall productivity growth. Second, because computers are a type of capital good, falling computer prices allow firms to accumulate more computing capital for every dollar of investment spending; the resulting increase in capital accumulation raises growth in all sectors that use computers as a factor of production. Third, the innovations in the computer industry may induce other industries to reconsider their own production methods, which in turn leads to productivity growth in those industries.

The big, open question is whether the computer industry will remain an engine of growth. Will Moore's law describe the future as well as it has described the past? Will the technological advances of the next decade be as profound as the Internet was during the 1990s? Stay tuned.⁵

8-3 From Growth Theory to Growth Empirics

So far in this chapter we have introduced exogenous technological progress into the Solow model to explain sustained growth in standards of living. We then used the theoretical framework as a lens through which to view some key issues facing policymakers. Let's now discuss what happens when the theory is asked to confront the facts.

Balanced Growth

According to the Solow model, technological progress causes the values of many variables to rise together in the steady state. This property, called *balanced growth*, does a good job of describing the long-run data for the U.S. economy.

Consider first output per worker Y/L and the capital stock per worker K/L . According to the Solow model, in the steady state, both of these variables grow at the rate of technological progress. United States data for the half century show that output per worker and the capital stock per worker have in fact grown at approximately the same rate—about 2 percent per year. To put it another way, the capital–output ratio has remained approximately constant over time.

⁵ For more on this topic, see the symposium on “Computers and Productivity” in the Fall 2000 issue of *The Journal of Economic Perspectives*. On the parallel between electricity and computers, see Paul A. David, “The Dynamo and the Computer: A Historical Perspective on the Modern Productivity Paradox,” *American Economic Review* 80, no. 2 (May 1990): 355–361.

Technological progress also affects factor prices. Problem 3(d) at the end of the chapter asks you to show that, in the steady state, the real wage grows at the rate of technological progress. The real rental price of capital, however, is constant over time. Again, these predictions hold true for the United States. Over the past 50 years, the real wage has increased about 2 percent per year; it has increased about the same amount as real GDP per worker. Yet the real rental price of capital (measured as real capital income divided by the capital stock) has remained about the same.

The Solow model's prediction about factor prices—and the success of this prediction—is especially noteworthy when contrasted with Karl Marx's theory of the development of capitalist economies. Marx predicted that the return to capital would decline over time and that this would lead to economic and political crises. Economic history has not supported Marx's prediction, which partly explains why we now study Solow's theory of growth rather than Marx's.

Convergence

If you travel around the world, you will see tremendous variations in living standards. The world's poor countries have average levels of income per person that are less than one-tenth the average levels in the world's rich countries. These differences in income are reflected in almost every measure of the quality of life—from the number of televisions and telephones per household to the infant mortality rate and life expectancy.

Much research has been devoted to the question of whether economies converge over time to one another. In particular, do economies that start off poor subsequently grow faster than economies that start off rich? If they do, then the world's poor economies will tend to catch up with the world's rich economies. This property of catch-up is called *convergence*.

To understand the study of convergence, consider an analogy. Imagine that you were to collect data on college students. At the end of their first year, some students have A averages, whereas others have C averages. Would you expect the A and the C students to converge over the remaining three years of college? The answer depends on why their first-year grades differed. If the differences arose because some students came from better high schools than others, then you might expect those who were initially disadvantaged to start catching up to their better-prepared peers. But if the differences arose because some students study more than others, you might expect the differences in grades to persist.

The Solow model predicts that much the same is true with nations: whether economies converge depends on why they differed in the first place. On the one hand, if two economies with the same steady state happened by historical accident to start off with different capital stocks, then we should expect them to converge. The economy with the smaller capital stock will naturally grow more quickly. (In a case study in Chapter 7, we applied this logic to explain rapid growth in Germany and Japan after World War II.) On the other hand, if two economies have different steady states, perhaps because the economies have different rates of saving, then we should not expect convergence. Instead, each economy will approach its own steady state.

Experience is consistent with this analysis. In samples of economies with similar cultures and policies, studies find that economies converge to one another at a rate of about 2 percent per year. That is, the gap between rich and poor economies closes by about 2 percent each year. An example is the economies of individual American states. For historical reasons, such as the Civil War of the 1860s, income levels varied greatly among states a century ago. Yet these differences have slowly disappeared over time.

In international data, a more complex picture emerges. When researchers examine only data on income per person, they find little evidence of convergence: countries that start off poor do not grow faster on average than countries that start off rich. This finding suggests that different countries have different steady states. If statistical techniques are used to control for some of the determinants of the steady state, such as saving rates, population growth rates, and educational attainment, then once again the data show convergence at a rate of about 2 percent per year. In other words, the economies of the world exhibit *conditional convergence*: they appear to be converging to their own steady states, which in turn are determined by saving, population growth, and education.⁶

Factor Accumulation Versus Production Efficiency

As a matter of accounting, international differences in income per person can be attributed to either (1) differences in the factors of production, such as the quantities of physical and human capital, or (2) differences in the efficiency with which economies use their factors of production. That is, a worker in a poor country may be poor because he lacks tools and skills or because the tools and skills he has are not being put to their best use. To describe this issue in terms of the Solow model, the question is whether the large gap between rich and poor is explained by differences in capital accumulation (including human capital) or differences in the production function.

Much research has attempted to estimate the relative importance of these two sources of income disparities. The exact answer varies from study to study, but both factor accumulation and production efficiency appear important. Moreover, a common finding is that they are positively correlated: nations with high levels of physical and human capital also tend to use those factors efficiently.⁷

There are several ways to interpret this positive correlation. One hypothesis is that an efficient economy may encourage capital accumulation. For example, a person in a well-functioning economy may have greater resources and incentive

⁶ Robert Barro and Xavier Sala-i-Martin, “Convergence Across States and Regions,” *Brookings Papers on Economic Activity*, no. 1 (1991): 107–182; and N. Gregory Mankiw, David Romer, and David N. Weil, “A Contribution to the Empirics of Economic Growth,” *Quarterly Journal of Economics* (May 1992): 407–437.

⁷ Robert E. Hall and Charles I. Jones, “Why Do Some Countries Produce So Much More Output per Worker Than Others?” *Quarterly Journal of Economics* 114 (February 1999): 83–116; and Peter J. Klenow and Andres Rodriguez-Clare, “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?” *NBER Macroeconomics Annual* (1997): 73–103.

to stay in school and accumulate human capital. Another hypothesis is that capital accumulation may induce greater efficiency. If there are positive externalities to physical and human capital, a possibility mentioned earlier in the chapter, then countries that save and invest more will appear to have better production functions (unless the research study accounts for these externalities, which is hard to do). Thus, greater production efficiency may cause greater factor accumulation, or the other way around.

A final hypothesis is that both factor accumulation and production efficiency are driven by a common third variable. Perhaps the common third variable is the quality of the nation's institutions, including the government's policymaking process. As one economist put it, when governments screw up, they screw up big time. Bad policies, such as high inflation, excessive budget deficits, widespread market interference, and rampant corruption, often go hand in hand. We should not be surprised that such economies both accumulate less capital and fail to use the capital they have as efficiently as they might.

8-4 Beyond the Solow Model: Endogenous Growth Theory

A chemist, a physicist, and an economist are all trapped on a desert island, trying to figure out how to open a can of food.

“Let’s heat the can over the fire until it explodes,” says the chemist.

“No, no,” says the physicist, “Let’s drop the can onto the rocks from the top of a high tree.”

“I have an idea,” says the economist. “First, we assume a can opener . . .”

This old joke takes aim at how economists use assumptions to simplify—and sometimes oversimplify—the problems they face. It is particularly apt when evaluating the theory of economic growth. One goal of growth theory is to explain the persistent rise in living standards that we observe in most parts of the world. The Solow growth model shows that such persistent growth must come from technological progress. But where does technological progress come from? In the Solow model, it is simply assumed!

To understand fully the process of economic growth, we need to go beyond the Solow model and develop models that explain technological progress. Models that do this often go by the label **endogenous growth theory**, because they reject the Solow model’s assumption of exogenous technological change. Although the field of endogenous growth theory is large and sometimes complex, here we get a quick sampling of this modern research.⁸

⁸ This section provides a brief introduction to the large and fascinating literature on endogenous growth theory. Early and important contributions to this literature include Paul M. Romer, “Increasing Returns and Long-Run Growth,” *Journal of Political Economy* 94 (October 1986): 1002–1037; and Robert E. Lucas, Jr., “On the Mechanics of Economic Development,” *Journal of Monetary Economics* 22 (1988): 3–42. The reader can learn more about this topic in the undergraduate textbook by Charles I. Jones, *Introduction to Economic Growth* (New York: Norton, 1998).

The Basic Model

To illustrate the idea behind endogenous growth theory, let's start with a particularly simple production function:

$$Y = AK,$$

where Y is output, K is the capital stock, and A is a constant measuring the amount of output produced for each unit of capital. Notice that this production function does not exhibit the property of diminishing returns to capital. One extra unit of capital produces A extra units of output, regardless of how much capital there is. This absence of diminishing returns to capital is the key difference between this endogenous growth model and the Solow model.

Now let's see how this production function relates to economic growth. As before, we assume a fraction s of income is saved and invested. We therefore describe capital accumulation with an equation similar to those we used previously:

$$\Delta K = sY - \delta K.$$

This equation states that the change in the capital stock (ΔK) equals investment (sY) minus depreciation (δK). Combining this equation with the $Y = AK$ production function, we obtain, after a bit of manipulation,

$$\Delta Y/Y = \Delta K/K = sA - \delta.$$

This equation shows what determines the growth rate of output $\Delta Y/Y$. Notice that, as long as $sA > \delta$, the economy's income grows forever, even without the assumption of exogenous technological progress.

Thus, a simple change in the production function can alter dramatically the predictions about economic growth. In the Solow model, saving leads to growth temporarily, but diminishing returns to capital eventually force the economy to approach a steady state in which growth depends only on exogenous technological progress. By contrast, in this endogenous growth model, saving and investment can lead to persistent growth.

But is it reasonable to abandon the assumption of diminishing returns to capital? The answer depends on how we interpret the variable K in the production function $Y = AK$. If we take the traditional view that K includes only the economy's stock of plants and equipment, then it is natural to assume diminishing returns. Giving 10 computers to each worker does not make the worker 10 times as productive as he or she is with one computer.

Advocates of endogenous growth theory, however, argue that the assumption of constant (rather than diminishing) returns to capital is more palatable if K is interpreted more broadly. Perhaps the best case for the endogenous growth model is to view knowledge as a type of capital. Clearly, knowledge is an important input into the economy's production—both its production of goods and services and its production of new knowledge. Compared to other forms of capital, however, it is less natural to assume that knowledge exhibits the property of diminishing returns. (Indeed, the increasing pace of scientific and technological innovation over the past few centuries has led some economists to argue that there are increasing

returns to knowledge.) If we accept the view that knowledge is a type of capital, then this endogenous growth model with its assumption of constant returns to capital becomes a more plausible description of long-run economic growth.

A Two-Sector Model

Although the $Y = AK$ model is the simplest example of endogenous growth, the theory has gone well beyond this. One line of research has tried to develop models with more than one sector of production in order to offer a better description of the forces that govern technological progress. To see what we might learn from such models, let's sketch out an example.

The economy has two sectors, which we can call manufacturing firms and research universities. Firms produce goods and services, which are used for consumption and investment in physical capital. Universities produce a factor of production called "knowledge," which is then freely used in both sectors. The economy is described by the production function for firms, the production function for universities, and the capital-accumulation equation:

$$\begin{aligned} Y &= F[K, (1 - u)EL] && \text{(production function in manufacturing firms),} \\ \Delta E &= g(u)E && \text{(production function in research universities),} \\ \Delta K &= sY - \delta K && \text{(capital accumulation),} \end{aligned}$$

where u is the fraction of the labor force in universities (and $1 - u$ is the fraction in manufacturing), E is the stock of knowledge (which in turn determines the efficiency of labor), and g is a function that shows how the growth in knowledge depends on the fraction of the labor force in universities. The rest of the notation is standard. As usual, the production function for the manufacturing firms is assumed to have constant returns to scale: if we double both the amount of physical capital (K) and the number of effective workers in manufacturing $[(1 - u)EL]$, we double the output of goods and services (Y).

This model is a cousin of the $Y = AK$ model. Most important, this economy exhibits constant (rather than diminishing) returns to capital, as long as capital is broadly defined to include knowledge. In particular, if we double both physical capital K and knowledge E , then we double the output of both sectors in the economy. As a result, like the $Y = AK$ model, this model can generate persistent growth without the assumption of exogenous shifts in the production function. Here persistent growth arises endogenously because the creation of knowledge in universities never slows down.

At the same time, however, this model is also a cousin of the Solow growth model. If u , the fraction of the labor force in universities, is held constant, then the efficiency of labor E grows at the constant rate $g(u)$. This result of constant growth in the efficiency of labor at rate g is precisely the assumption made in the Solow model with technological progress. Moreover, the rest of the model—the manufacturing production function and the capital-accumulation equation—also resembles the rest of the Solow model. As a result, for any given value of u , this endogenous growth model works just like the Solow model.

There are two key decision variables in this model. As in the Solow model, the fraction of output used for saving and investment, s , determines the steady-state stock of physical capital. In addition, the fraction of labor in universities, u , determines the growth in the stock of knowledge. Both s and u affect the level of income, although only u affects the steady-state growth rate of income. Thus, this model of endogenous growth takes a small step in the direction of showing which societal decisions determine the rate of technological change.

The Microeconomics of Research and Development

The two-sector endogenous growth model just presented takes us closer to understanding technological progress, but it still tells only a rudimentary story about the creation of knowledge. If one thinks about the process of research and development for even a moment, three facts become apparent. First, although knowledge is largely a public good (that is, a good freely available to everyone), much research is done in firms that are driven by the profit motive. Second, research is profitable because innovations give firms temporary monopolies, either because of the patent system or because there is an advantage to being the first firm on the market with a new product. Third, when one firm innovates, other firms build on that innovation to produce the next generation of innovations. These (essentially microeconomic) facts are not easily connected with the (essentially macroeconomic) growth models we have discussed so far.

Some endogenous growth models try to incorporate these facts about research and development. Doing this requires modeling the decisions that firms face as they engage in research and modeling the interactions among firms that have some degree of monopoly power over their innovations. Going into more detail about these models is beyond the scope of this book. But it should be clear already that one virtue of these endogenous growth models is that they offer a more complete description of the process of technological innovation.

One question these models are designed to address is whether, from the standpoint of society as a whole, private profit-maximizing firms tend to engage in too little or too much research. In other words, is the social return to research (which is what society cares about) greater or smaller than the private return (which is what motivates individual firms)? It turns out that, as a theoretical matter, there are effects in both directions. On the one hand, when a firm creates a new technology, it makes other firms better off by giving them a base of knowledge on which to build future research. As Isaac Newton famously remarked, “If I have seen farther than others, it is because I was standing on the shoulders of giants.” On the other hand, when one firm invests in research, it can also make other firms worse off by merely being first to discover a technology that another firm would have invented. This duplication of research effort has been called the “stepping on toes” effect. Whether firms left to their own devices do too little or too much research depends on whether the positive “standing on shoulders” externality or the negative “stepping on toes” externality is more prevalent.

Although theory alone is ambiguous about the optimality of research effort, the empirical work in this area is usually less so. Many studies have suggested the “standing on shoulders” externality is important and, as a result, the social return to research is large—often in excess of 40 percent per year. This is an impressive rate of return, especially when compared to the return to physical capital, which we earlier estimated to be about 8 percent per year. In the judgment of some economists, this finding justifies substantial government subsidies to research.⁹

8-5 Conclusion

Long-run economic growth is the single most important determinant of the economic well-being of a nation’s citizens. Everything else that macroeconomists study—unemployment, inflation, trade deficits, and so on—pales in comparison.

Fortunately, economists know quite a lot about the forces that govern economic growth. The Solow growth model and the more recent endogenous growth models show how saving, population growth, and technological progress interact in determining the level and growth of a nation’s standard of living. Although these theories offer no magic pill to ensure an economy achieves rapid growth, they do offer much insight, and they provide the intellectual framework for much of the debate over public policy.

Summary

1. In the steady state of the Solow growth model, the growth rate of income per person is determined solely by the exogenous rate of technological progress.
2. In the Solow model with population growth and technological progress, the Golden Rule (consumption-maximizing) steady state is characterized by equality between the net marginal product of capital ($MPK - \delta$) and the steady-state growth rate ($n + g$). By contrast, in the U.S. economy, the net marginal product of capital is well in excess of the growth rate, indicating that the U.S. economy has much less capital than in the Golden Rule steady state.
3. Policymakers in the United States and other countries often claim that their nations should devote a larger percentage of their output to saving and investment. Increased public saving and tax incentives for private saving are two ways to encourage capital accumulation.
4. In the early 1970s, the rate of growth fell substantially in most industrialized countries. The cause of this slowdown is not well understood. In the mid-1990s, the rate of growth increased, most likely because of advances in information technology.

⁹ For an overview of the empirical literature on the effects of research, see Zvi Griliches, “The Search for R&D Spillovers,” *Scandinavian Journal of Economics* 94 (1991): 29–47.

5. Many empirical studies have examined to what extent the Solow model can help explain long-run economic growth. The model can explain much of what we see in the data, such as balanced growth and conditional convergence. Recent studies have also found that international variation in standards of living is attributable to a combination of capital accumulation and the efficiency with which capital is used.
6. Modern theories of endogenous growth attempt to explain the rate of technological progress, which the Solow model takes as exogenous. These models try to explain the decisions that determine the creation of knowledge through research and development.

KEY CONCEPTS

Efficiency of labor

Labor-augmenting
technological progress

Endogenous growth theory

QUESTIONS FOR REVIEW

1. In the Solow model, what determines the steady-state rate of growth of income per worker?
2. What data would you need to determine whether an economy has more or less capital than in the Golden Rule steady state?
3. How can policymakers influence a nation's saving rate?
4. What has happened to the rate of productivity growth over the past 40 years? How might you explain this phenomenon?
5. In the steady state of the Solow model, at what rate does output per person grow? At what rate does capital per person grow? How does this compare with U.S. experience?
6. How does endogenous growth theory explain persistent growth without the assumption of exogenous technological progress? How does this differ from the Solow model?

PROBLEMS AND APPLICATIONS

1. An economy described by the Solow growth model has the following production function:

$$y = \sqrt{k}.$$

- a. Solve for the steady-state value of y as a function of s , n , g , and δ .
- b. A developed country has a saving rate of 28 percent and a population growth rate of 1 percent per year. A less-developed country has a saving rate of 10 percent and a population growth rate of 4 percent per year. In both countries, $g = 0.02$ and $\delta = 0.04$. Find the steady-state value of y for each country.
- c. What policies might the less-developed country pursue to raise its level of income?
2. In the United States, the capital share of GDP is about 30 percent; the average growth in output is about 3 percent per year; the depreciation rate is about 4 percent per year; and the capital-output ratio is about 2.5. Suppose that the production function is Cobb–Douglas, so that the capital

share in output is constant, and that the United States has been in a steady state. (For a discussion of the Cobb–Douglas production function, see the appendix to Chapter 3.)

- a. What must the saving rate be in the initial steady state? [*Hint*: Use the steady-state relationship, $sy = (\delta + n + g)k$.]
 - b. What is the marginal product of capital in the initial steady state?
 - c. Suppose that public policy raises the saving rate so that the economy reaches the Golden Rule level of capital. What will the marginal product of capital be at the Golden Rule steady state? Compare the marginal product at the Golden Rule steady state to the marginal product in the initial steady state. Explain.
 - d. What will the capital–output ratio be at the Golden Rule steady state? [*Hint*: For the Cobb–Douglas production function, the capital–output ratio is related to the marginal product of capital.]
 - e. What must the saving rate be to reach the Golden Rule steady state?
3. Prove each of the following statements about the steady state with population growth and technological progress.
- a. The capital–output ratio is constant.
 - b. Capital and labor each earn a constant share of an economy’s income. [*Hint*: Recall the definition $MPK = f(k + 1) - f(k)$.]
 - c. Total capital income and total labor income both grow at the rate of population growth plus the rate of technological progress, $n + g$.
 - d. The real rental price of capital is constant, and the real wage grows at the rate of technological progress g . [*Hint*: The real rental price of capital equals total capital income divided by the capital stock, and the real wage equals total labor income divided by the labor force.]
4. The amount of education the typical person receives varies substantially among countries. Suppose you were to compare a country with a

highly educated labor force and a country with a less educated labor force. Assume that education affects only the level of the efficiency of labor. Also assume that the countries are otherwise the same: they have the same saving rate, the same depreciation rate, the same population growth rate, and the same rate of technological progress. Both countries are described by the Solow model and are in their steady states. What would you predict for the following variables?

- a. The rate of growth of total income.
 - b. The level of income per worker.
 - c. The real rental price of capital.
 - d. The real wage.
5. This question asks you to analyze in more detail the two-sector endogenous growth model presented in the text.
- a. Rewrite the production function for manufactured goods in terms of output per effective worker and capital per effective worker.
 - b. In this economy, what is break-even investment (the amount of investment needed to keep capital per effective worker constant)?
 - c. Write down the equation of motion for k , which shows Δk as saving minus break-even investment. Use this equation to draw a graph showing the determination of steady-state k . [*Hint*: This graph will look much like those we used to analyze the Solow model.]
 - d. In this economy, what is the steady-state growth rate of output per worker Y/L ? How do the saving rate s and the fraction of the labor force in universities u affect this steady-state growth rate?
 - e. Using your graph, show the impact of an increase in u . [*Hint*: This change affects both curves.] Describe both the immediate and the steady-state effects.
 - f. Based on your analysis, is an increase in u an unambiguously good thing for the economy? Explain.

APPENDIX

Accounting for the Sources of Economic Growth

Real GDP in the United States has grown an average of 3 percent per year over the past 40 years. What explains this growth? In Chapter 3 we linked the output of the economy to the factors of production—capital and labor—and to the production technology. Here we develop a technique called *growth accounting* that divides the growth in output into three different sources: increases in capital, increases in labor, and advances in technology. This breakdown provides us with a measure of the rate of technological change.

Increases in the Factors of Production

We first examine how increases in the factors of production contribute to increases in output. To do this, we start by assuming there is no technological change, so the production function relating output Y to capital K and labor L is constant over time:

$$Y = F(K, L).$$

In this case, the amount of output changes only because the amount of capital or labor changes.

Increases in Capital First, consider changes in capital. If the amount of capital increases by ΔK units, by how much does the amount of output increase? To answer this question, we need to recall the definition of the marginal product of capital MPK :

$$MPK = F(K + 1, L) - F(K, L).$$

The marginal product of capital tells us how much output increases when capital increases by 1 unit. Therefore, when capital increases by ΔK units, output increases by approximately $MPK \times \Delta K$.¹⁰

For example, suppose that the marginal product of capital is $1/5$; that is, an additional unit of capital increases the amount of output produced by one-fifth of a

¹⁰ Note the word “approximately” here. This answer is only an approximation because the marginal product of capital varies: it falls as the amount of capital increases. An exact answer would take into account the fact that each unit of capital has a different marginal product. If the change in K is not too large, however, the approximation of a constant marginal product is very accurate.

unit. If we increase the amount of capital by 10 units, we can compute the amount of additional output as follows:

$$\begin{aligned}\Delta Y &= MPK \times \Delta K \\ &= 1/5 \frac{\text{Units of Output}}{\text{Unit of Capital}} \times 10 \text{ Units of Capital} \\ &= 2 \text{ Units of Output.}\end{aligned}$$

By increasing capital 10 units, we obtain 2 more units of output. Thus, we use the marginal product of capital to convert changes in capital into changes in output.

Increases in Labor Next, consider changes in labor. If the amount of labor increases by ΔL units, by how much does output increase? We answer this question the same way we answered the question about capital. The marginal product of labor MPL tells us how much output changes when labor increases by 1 unit—that is,

$$MPL = F(K, L + 1) - F(K, L).$$

Therefore, when the amount of labor increases by ΔL units, output increases by approximately $MPL \times \Delta L$.

For example, suppose that the marginal product of labor is 2; that is, an additional unit of labor increases the amount of output produced by 2 units. If we increase the amount of labor by 10 units, we can compute the amount of additional output as follows:

$$\begin{aligned}\Delta Y &= MPL \times \Delta L \\ &= 2 \frac{\text{Units of Output}}{\text{Unit of Labor}} \times 10 \text{ Units of Labor} \\ &= 20 \text{ Units of Output.}\end{aligned}$$

By increasing labor 10 units, we obtain 20 more units of output. Thus, we use the marginal product of labor to convert changes in labor into changes in output.

Increases in Capital and Labor Finally, let's consider the more realistic case in which both factors of production change. Suppose that the amount of capital increases by ΔK and the amount of labor increases by ΔL . The increase in output then comes from two sources: more capital and more labor. We can divide this increase into the two sources using the marginal products of the two inputs:

$$\Delta Y = (MPK \times \Delta K) + (MPL \times \Delta L).$$

The first term in parentheses is the increase in output resulting from the increase in capital, and the second term in parentheses is the increase in output resulting from the increase in labor. This equation shows us how to attribute growth to each factor of production.

We now want to convert this last equation into a form that is easier to interpret and apply to the available data. First, with some algebraic rearrangement, the equation becomes¹¹

$$\frac{\Delta Y}{Y} = \left(\frac{MPK \times K}{Y} \right) \frac{\Delta K}{K} + \left(\frac{MPL \times L}{Y} \right) \frac{\Delta L}{L}.$$

This form of the equation relates the growth rate of output, $\Delta Y/Y$, to the growth rate of capital, $\Delta K/K$, and the growth rate of labor, $\Delta L/L$.

Next, we need to find some way to measure the terms in parentheses in the last equation. In Chapter 3 we showed that the marginal product of capital equals its real rental price. Therefore, $MPK \times K$ is the total return to capital, and $(MPK \times K)/Y$ is capital's share of output. Similarly, the marginal product of labor equals the real wage. Therefore, $MPL \times L$ is the total compensation that labor receives, and $(MPL \times L)/Y$ is labor's share of output. Under the assumption that the production function has constant returns to scale, Euler's theorem (which we discussed in Chapter 3) tells us that these two shares sum to 1. In this case, we can write

$$\frac{\Delta Y}{Y} = \alpha \frac{\Delta K}{K} + (1 - \alpha) \frac{\Delta L}{L}.$$

where α is capital's share and $(1 - \alpha)$ is labor's share.

This last equation gives us a simple formula for showing how changes in inputs lead to changes in output. In particular, we must weight the growth rates of the inputs by the factor shares. As we discussed in the appendix to Chapter 3, capital's share in the United States is about 30 percent, that is, $\alpha = 0.30$. Therefore, a 10-percent increase in the amount of capital ($\Delta K/K = 0.10$) leads to a 3-percent increase in the amount of output ($\Delta Y/Y = 0.03$). Similarly, a 10-percent increase in the amount of labor ($\Delta L/L = 0.10$) leads to a 7-percent increase in the amount of output ($\Delta Y/Y = 0.07$).

Technological Progress

So far in our analysis of the sources of growth, we have been assuming that the production function does not change over time. In practice, of course, technological progress improves the production function. For any given amount of inputs, we get more output today than we did in the past. We now extend the analysis to allow for technological progress.

¹¹ *Mathematical note:* To see that this is equivalent to the previous equation, note that we can multiply both sides of this equation by Y and thereby cancel Y from three places in which it appears. We can cancel the K in the top and bottom of the first term on the right-hand side and the L in the top and bottom of the second term on the right-hand side. These algebraic manipulations turn this equation into the previous one.

We include the effects of the changing technology by writing the production function as

$$Y = AF(K, L),$$

where A is a measure of the current level of technology called *total factor productivity*. Output now increases not only because of increases in capital and labor but also because of increases in total factor productivity. If total factor productivity increases by 1 percent and if the inputs are unchanged, then output increases by 1 percent.

Allowing for a changing technology adds another term to our equation accounting for economic growth:

$$\frac{\Delta Y}{Y} = \alpha \frac{\Delta K}{K} + (1 - \alpha) \frac{\Delta L}{L} + \frac{\Delta A}{A}$$

Growth in Contribution Contribution Growth in Total
Output of Capital of Labor Factor Productivity

This is the key equation of growth accounting. It identifies and allows us to measure the three sources of growth: changes in the amount of capital, changes in the amount of labor, and changes in total factor productivity.

Because total factor productivity is not observable directly, it is measured indirectly. We have data on the growth in output, capital, and labor; we also have data on capital's share of output. From these data and the growth-accounting equation, we can compute the growth in total factor productivity to make sure that everything adds up:

$$\frac{\Delta A}{A} = \frac{\Delta Y}{Y} - \alpha \frac{\Delta K}{K} - (1 - \alpha) \frac{\Delta L}{L}.$$

$\Delta A/A$ is the change in output that cannot be explained by changes in inputs. Thus, the growth in total factor productivity is computed as a residual—that is, as the amount of output growth that remains after we have accounted for the determinants of growth that we can measure. Indeed, $\Delta A/A$ is sometimes called the *Solow residual*, after Robert Solow, who first showed how to compute it.¹²

Total factor productivity can change for many reasons. Changes most often arise because of increased knowledge about production methods, and the Solow residual is often used as a measure of technological progress. Yet other

¹² Robert M. Solow, "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics* 39 (1957): 312–320. It is natural to ask how growth in labor efficiency E relates to growth in total factor productivity. One can show that $\Delta A/A = (1 - \alpha)\Delta E/E$, where α is capital's share. Thus, technological change as measured by growth in the efficiency of labor is proportional to technological change as measured by the Solow residual.

factors, such as education and government regulation, can affect total factor productivity as well. For example, if higher public spending raises the quality of education, then workers may become more productive and output may rise, which implies higher total factor productivity. As another example, if government regulations require firms to purchase capital to reduce pollution or increase worker safety, then the capital stock may rise without any increase in measured output, which implies lower total factor productivity. *Total factor productivity captures anything that changes the relation between measured inputs and measured output.*

The Sources of Growth in the United States

Having learned how to measure the sources of economic growth, we now look at the data. Table 8-3 uses U.S. data to measure the contributions of the three sources of growth between 1950 and 1999.

This table shows that real GDP has grown an average of 3.6 percent per year since 1950. Of this 3.6 percent, 1.2 percent is attributable to increases in the capital stock, 1.3 percent to increases in the labor input, and 1.1 percent to increases in total factor productivity. These data show that increases in capital, labor, and productivity have contributed almost equally to economic growth in the United States.

Table 8-3 also shows that the growth in total factor productivity slowed substantially around 1970. In a previous case study in this chapter, we discussed some hypotheses to explain this productivity slowdown.

table 8-3

Accounting for Economic Growth in the United States

Years	Output Growth $\Delta Y/Y$	SOURCE OF GROWTH		
		Capital $\alpha \Delta K/K$	Labor $(1 - \alpha) \Delta K/K$	Total Factor Productivity $\Delta A/A$
		(average percentage increase per year)		
1950–1999	3.6	1.2	1.3	1.1
1950–1960	3.3	1.0	1.0	1.3
1960–1970	4.4	1.4	1.2	1.8
1970–1980	3.6	1.4	1.2	1.0
1980–1990	3.4	1.2	1.6	0.6
1990–1999	3.7	1.2	1.6	0.9

Source: U.S. Department of Commerce, U.S. Department of Labor, and the author's calculations. The parameter α is set to equal 0.3.

CASE STUDY

Growth in the East Asian Tigers

Perhaps the most spectacular growth experiences in recent history have been those of the “Tigers” of East Asia: Hong Kong, Singapore, South Korea, and Taiwan. From 1966 to 1990, while real income per person was growing about 2 percent per year in the United States, it grew more than 7 percent per year in each of these countries. In the course of a single generation, real income per person increased fivefold, moving the Tigers from among the world’s poorest countries to among the richest. (In the late 1990s, a period of pronounced financial turmoil tarnished the reputation of some of these economies. But this short-run problem, which we examine in a case study in Chapter 12, doesn’t come close to reversing the spectacular long-run growth performance that the Asian Tigers have experienced.)

What accounts for these growth miracles? Some commentators have argued that the success of these four countries is hard to reconcile with basic growth theory, such as the Solow growth model, which has technology growing at a constant, exogenous rate. They have suggested that these countries’ rapid growth is explained by their ability to imitate foreign technologies. By adopting technology developed abroad, the argument goes, these countries managed to improve their production functions substantially in a relatively short period of time. If this argument is correct, these countries should have experienced unusually rapid growth in total factor productivity.

One recent study shed light on this issue by examining in detail the data from these four countries. The study found that their exceptional growth can be traced to large increases in measured factor inputs: increases in labor-force participation, increases in the capital stock, and increases in educational attainment. In South Korea, for example, the investment–GDP ratio rose from about 5 percent in the 1950s to about 30 percent in the 1980s; the percentage of the working population with at least a high-school education went from 26 percent in 1966 to 75 percent in 1991.

Once we account for growth in labor, capital, and human capital, little of the growth in output is left to explain. None of these four countries experienced unusually rapid growth in total factor productivity. Indeed, the average growth in total factor productivity in the East Asian Tigers was almost exactly the same as in the United States. Thus, although these countries’ rapid growth has been truly impressive, it is easy to explain using the tools of basic growth theory.¹³

¹³ Alwyn Young, “The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience,” *Quarterly Journal of Economics* 101 (August 1995): 641–680.

MORE PROBLEMS AND APPLICATIONS

1. In the economy of Solovia, the owners of capital get two-thirds of national income, and the workers receive one-third.
 - a. The men of Solovia stay at home performing household chores, while the women work in factories. If some of the men started working outside the home so that the labor force increased by 5 percent, what would happen to the measured output of the economy? Does labor productivity—defined as output per worker—increase, decrease, or stay the same? Does total factor productivity increase, decrease, or stay the same?
 - b. In year 1, the capital stock was 6, the labor input was 3, and output was 12. In year 2, the capital stock was 7, the labor input was 4, and output was 14. What happened to total factor productivity between the two years?
2. Labor productivity is defined as Y/L , the amount of output divided by the amount of labor input. Start with the growth-accounting equation and show that the growth in labor productivity depends on growth in total factor productivity and

growth in the capital-labor ratio. In particular, show that

$$\frac{\Delta(Y/L)}{Y/L} = \frac{\Delta A}{A} + \alpha \frac{\Delta(K/L)}{K/L}.$$

(*Hint:* You may find the following mathematical trick helpful. If $z = wx$, then the growth rate of z is approximately the growth rate of w plus the growth rate of x . That is,

$$\Delta z/z \approx \Delta w/w + \Delta x/x.)$$

3. Suppose an economy described by the Solow model is in a steady state with population growth n of 1.0 percent per year and technological progress g of 2.0 percent per year. Total output and total capital grow at 3.0 percent per year. Suppose further that the capital share of output is 0.3. If you used the growth-accounting equation to divide output growth into three sources—capital, labor, and total factor productivity—how much would you attribute to each source? Compare your results to the figures we found for the United States in Table 8-3.

part IV

Business Cycle Theory: The Economy in the Short Run

Introduction to Economic Fluctuations

The modern world regards business cycles much as the ancient Egyptians regarded the overflowing of the Nile. The phenomenon recurs at intervals, it is of great importance to everyone, and natural causes of it are not in sight.

— John Bates Clark, 1898


Economic fluctuations present a recurring problem for economists and policy-makers. This problem is illustrated in Figure 9-1, which shows growth in real GDP for the U.S. economy. As you can see, although the economy experiences long-run growth that averages about 3.5 percent per year, this growth is not at all steady. Recessions—periods of falling incomes and rising unemployment—are frequent. In the recession of 1990, for instance, real GDP fell 2.2 percent from its peak to its trough, and the unemployment rate rose to 7.7 percent. During recessions, not only are more people unemployed, but those who are employed have shorter workweeks, as more workers have to accept part-time jobs and fewer workers have the opportunity to work overtime. When recessions end and the economy enters a boom, these effects work in reverse: incomes rise, unemployment falls, and workweeks expand.

Economists call these short-run fluctuations in output and employment the *business cycle*. Although this term suggests that economic fluctuations are regular and predictable, they are not. Recessions are as irregular as they are common. Sometimes they are close together, such as the recessions of 1980 and 1982. Sometimes they are far apart, such as the recessions of 1982 and 1990.

In Parts II and III of this book, we developed theories to explain how the economy behaves in the long run. Those theories were based on the classical dichotomy—the premise that real variables, such as output and employment, are not affected by what happens to nominal variables, such as the money supply and the price level. Although classical theories are useful for explaining long-run trends, including the economic growth we observe from decade to decade, most economists believe that the classical dichotomy does not hold in the short run and, therefore, that classical theories cannot explain year-to-year fluctuations in output and employment.

Here, in Part IV, we see how economists explain these short-run fluctuations. This chapter begins our analysis by discussing the key differences between the long run and the short run and by introducing the model of aggregate supply

figure 9-1



Real GDP Growth in the United States Growth in real GDP averages about 3.5 percent per year, as indicated by the orange line, but there are substantial fluctuations around this average. Recessions are periods when the production of goods and services is declining, represented here by negative growth in real GDP.

Source: U.S. Department of Commerce.

and aggregate demand. With this model we can show how shocks to the economy lead to short-run fluctuations in output and employment.

Just as Egypt now controls the flooding of the Nile Valley with the Aswan Dam, modern society tries to control the business cycle with appropriate economic policies. The model we develop over the next several chapters shows how monetary and fiscal policies influence the business cycle. We will see that these policies can potentially stabilize the economy or, if poorly conducted, make the problem of economic instability even worse.

9-1 Time Horizons in Macroeconomics

Before we start building a model of short-run economic fluctuations, let's step back and ask a fundamental question: Why do economists need different models for different time horizons? Why can't we stop the course here and be content with the classical models developed in Chapters 3 through 8? The answer, as this book has consistently reminded its reader, is that classical macroeconomic theory applies to the long run but not to the short run. But why is this so?

How the Short Run and Long Run Differ

Most macroeconomists believe that the key difference between the short run and the long run is the behavior of prices. *In the long run, prices are flexible and can respond to changes in supply or demand. In the short run, many prices are “sticky” at some predetermined level.* Because prices behave differently in the short run than in the long run, economic policies have different effects over different time horizons.

To see how the short run and the long run differ, consider the effects of a change in monetary policy. Suppose that the Federal Reserve suddenly reduced the money supply by 5 percent. According to the classical model, which almost all economists agree describes the economy in the long run, the money supply affects nominal variables—variables measured in terms of money—but not real variables. In the long run, a 5-percent reduction in the money supply lowers all prices (including nominal wages) by 5 percent whereas all real variables remain the same. Thus, in the long run, changes in the money supply do not cause fluctuations in output or employment.

In the short run, however, many prices do not respond to changes in monetary policy. A reduction in the money supply does not immediately cause all firms to cut the wages they pay, all stores to change the price tags on their goods, all mail-order firms to issue new catalogs, and all restaurants to print new menus. Instead, there is little immediate change in many prices; that is, many prices are sticky. This short-run price stickiness implies that the short-run impact of a change in the money supply is not the same as the long-run impact.

A model of economic fluctuations must take into account this short-run price stickiness. We will see that the failure of prices to adjust quickly and completely means that, in the short run, output and employment must do some of the adjusting instead. In other words, during the time horizon over which prices are sticky, the classical dichotomy no longer holds: nominal variables can influence real variables, and the economy can deviate from the equilibrium predicted by the classical model.

CASE STUDY

The Puzzle of Sticky Magazine Prices

How sticky are prices? The answer to this question depends on what price we consider. Some commodities, such as wheat, soybeans, and pork bellies, are traded on organized exchanges, and their prices change every minute. No one would call these prices sticky. Yet the prices of most goods and services change much less frequently. One survey found that 39 percent of firms change their prices once a year, and another 10 percent change their prices less than once a year.¹

The reasons for price stickiness are not always apparent. Consider, for example, the market for magazines. A study has documented that magazines change their newsstand prices very infrequently. The typical magazine allows inflation to erode

¹ Alan S. Blinder, “On Sticky Prices: Academic Theories Meet the Real World,” in N. G. Mankiw, ed., *Monetary Policy* (Chicago: University of Chicago Press, 1994): 117–154. A case study in Chapter 19 discusses this survey in more detail.

its real price by about 25 percent before it raises its nominal price. When inflation is 4 percent per year, the typical magazine changes its price about every six years.²

Why do magazines keep their prices the same for so long? Economists do not have a definitive answer. The question is puzzling because it would seem that for magazines, the cost of a price change is small. To change prices, a mail-order firm must issue a new catalog and a restaurant must print a new menu, but a magazine publisher can simply print a new price on the cover of the next issue. Perhaps the cost to the publisher of charging the wrong price is also small. Or maybe customers would find it inconvenient if the price of their favorite magazine changed every month.

As the magazine example shows, explaining at the microeconomic level why prices are sticky is sometimes hard. The cause of price stickiness is, therefore, an active area of research, which we discuss more fully in Chapter 19. In this chapter, however, we simply assume that prices are sticky so we can start developing the link between sticky prices and the business cycle. Although not yet fully explained, short-run price stickiness is widely believed to be crucial for understanding short-run economic fluctuations.

The Model of Aggregate Supply and Aggregate Demand

How does introducing sticky prices change our view of how the economy works? We can answer this question by considering economists' two favorite words—supply and demand.

In classical macroeconomic theory, the amount of output depends on the economy's ability to *supply* goods and services, which in turn depends on the supplies of capital and labor and on the available production technology. This is the essence of the basic classical model in Chapter 3, as well as of the Solow growth model in Chapters 7 and 8. Flexible prices are a crucial assumption of classical theory. The theory posits, sometimes implicitly, that prices adjust to ensure that the quantity of output demanded equals the quantity supplied.

The economy works quite differently when prices are sticky. In this case, as we will see, output also depends on the *demand* for goods and services. Demand, in turn, is influenced by monetary policy, fiscal policy, and various other factors. Because monetary and fiscal policy can influence the economy's output over the time horizon when prices are sticky, price stickiness provides a rationale for why these policies may be useful in stabilizing the economy in the short run.

In the rest of this chapter, we develop a model that makes these ideas more precise. The model of supply and demand, which we used in Chapter 1 to discuss the market for pizza, offers some of the most fundamental insights in economics. This model shows how the supply and demand for any good jointly determine the

² Stephen G. Cecchetti, "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," *Journal of Econometrics* 31 (1986): 255–274.

good's price and the quantity sold, and how shifts in supply and demand affect the price and quantity. In the rest of this chapter, we introduce the “economy-size” version of this model—the *model of aggregate supply and aggregate demand*. This macroeconomic model allows us to study how the aggregate price level and the quantity of aggregate output are determined. It also provides a way to contrast how the economy behaves in the long run and how it behaves in the short run.

Although the model of aggregate supply and aggregate demand resembles the model of supply and demand for a single good, the analogy is not exact. The model of supply and demand for a single good considers only one good within a large economy. By contrast, as we will see in the coming chapters, the model of aggregate supply and aggregate demand is a sophisticated model that incorporates the interactions among many markets.

9-2 Aggregate Demand

Aggregate demand (AD) is the relationship between the quantity of output demanded and the aggregate price level. In other words, the aggregate demand curve tells us the quantity of goods and services people want to buy at any given level of prices. We examine the theory of aggregate demand in detail in Chapters 10 through 12. Here we use the quantity theory of money to provide a simple, although incomplete, derivation of the aggregate demand curve.

The Quantity Equation as Aggregate Demand

Recall from Chapter 4 that the quantity theory says that

$$MV = PY,$$

where M is the money supply, V is the velocity of money, P is the price level, and Y is the amount of output. If the velocity of money is constant, then this equation states that the money supply determines the nominal value of output, which in turn is the product of the price level and the amount of output.


You might recall that the quantity equation can be rewritten in terms of the supply and demand for real money balances:

$$M/P = (M/P)^d = kY,$$

where $k = 1/V$ is a parameter determining how much money people want to hold for every dollar of income. In this form, the quantity equation states that the supply of real money balances M/P equals the demand $(M/P)^d$ and that the demand is proportional to output Y . The velocity of money V is the “flip side” of the money demand parameter k .

For any fixed money supply and velocity, the quantity equation yields a negative relationship between the price level P and output Y . Figure 9-2 graphs the combinations of P and Y that satisfy the quantity equation holding M and V constant. This downward-sloping curve is called the aggregate demand curve.

figure 9-2

**The Aggregate Demand Curve**

The aggregate demand curve AD shows the relationship between the price level P and the quantity of goods and services demanded Y . It is drawn for a given value of the money supply M . The aggregate demand curve slopes downward: the higher the price level P , the lower the level of real balances M/P , and therefore the lower the quantity of goods and services demanded Y .

Why the Aggregate Demand Curve Slopes Downward

As a strictly mathematical matter, the quantity equation explains the downward slope of the aggregate demand curve very simply. The money supply M and the velocity of money V determine the nominal value of output PY . Once PY is fixed, if P goes up, Y must go down.

What is the economics that lies behind this mathematical relationship? For a complete answer, we have to wait for a couple of chapters. For now, however, consider the following logic: Because we have assumed the velocity of money is fixed, the money supply determines the dollar value of all transactions in the economy. (This conclusion should be familiar from Chapter 4.) If the price level rises, each transaction requires more dollars, so the number of transactions and thus the quantity of goods and services purchased must fall.


We can also explain the downward slope of the aggregate demand curve by thinking about the supply and demand for real money balances. If output is higher, people engage in more transactions and need higher real balances M/P . For a fixed money supply M , higher real balances imply a lower price level. Conversely, if the price level is lower, real money balances are higher; the higher level of real balances allows a greater volume of transactions, which means a greater quantity of output is demanded.

Shifts in the Aggregate Demand Curve

The aggregate demand curve is drawn for a fixed value of the money supply. In other words, it tells us the possible combinations of P and Y for a given value of M . If the Fed changes the money supply, then the possible combinations of P and Y change, which means the aggregate demand curve shifts.


For example, consider what happens if the Fed reduces the money supply. The quantity equation, $MV = PY$, tells us that the reduction in the money supply leads to a proportionate reduction in the nominal value of output PY . For any given price level, the amount of output is lower, and for any given amount of output, the price level is lower. As in Figure 9-3(a), the aggregate demand curve relating P and Y shifts inward.

figure 9-3



Shifts in the Aggregate Demand Curve

Changes in the money supply shift the aggregate demand curve. In panel (a), a decrease in the money supply M reduces the nominal value of output PY . For any given price level P , output Y is lower. Thus, a decrease in the money supply shifts the aggregate demand curve inward from AD_1 to AD_2 . In panel (b), an increase in the money supply M raises the nominal value of output PY . For any given price level P , output Y is higher. Thus, an increase in the money supply shifts the aggregate demand curve outward from AD_1 to AD_2 .



The opposite occurs if the Fed increases the money supply. The quantity equation tells us that an increase in M leads to an increase in PY . For any given price level, the amount of output is higher, and for any given amount of output, the price level is higher. As shown in Figure 9-3(b), the aggregate demand curve shifts outward.

Fluctuations in the money supply are not the only source of fluctuations in aggregate demand. Even if the money supply is held constant, the aggregate demand curve shifts if some event causes a change in the velocity of money. Over the next three chapters, we consider many possible reasons for shifts in the aggregate demand curve.

9-3 Aggregate Supply

By itself, the aggregate demand curve does not tell us the price level or the amount of output; it merely gives a relationship between these two variables. To accompany the aggregate demand curve, we need another relationship between P and Y that crosses the aggregate demand curve—an aggregate supply curve. The aggregate demand and aggregate supply curves together pin down the economy's price level and quantity of output.

Aggregate supply (AS) is the relationship between the quantity of goods and services supplied and the price level. Because the firms that supply goods and services have flexible prices in the long run but sticky prices in the short run, the aggregate supply relationship depends on the time horizon. We need to discuss two different aggregate supply curves: the long-run aggregate supply curve $LRAS$ and the short-run aggregate supply curve $SRAS$. We also need to discuss how the economy makes the transition from the short run to the long run.

The Long Run: The Vertical Aggregate Supply Curve


Because the classical model describes how the economy behaves in the long run, we derive the long-run aggregate supply curve from the classical model. Recall from Chapter 3 that the amount of output produced depends on the fixed amounts of capital and labor and on the available technology. To show this, we write

$$\begin{aligned} Y &= F(\bar{K}, \bar{L}) \\ &= \bar{Y}. \end{aligned}$$

According to the classical model, output does not depend on the price level. To show that output is the same for all price levels, we draw a vertical aggregate supply curve, as in Figure 9-4. The intersection of the aggregate demand curve with this vertical aggregate supply curve determines the price level.

If the aggregate supply curve is vertical, then changes in aggregate demand affect prices but not output. For example, if the money supply falls, the aggregate

figure 9-4




The Long-Run Aggregate Supply Curve In the long run, the level of output is determined by the amounts of capital and labor and by the available technology; it does not depend on the price level. The long-run aggregate supply curve, *LRAS*, is vertical.

demand curve shifts downward, as in Figure 9-5. The economy moves from the old intersection of aggregate supply and aggregate demand, point A, to the new intersection, point B. The shift in aggregate demand affects only prices.

The vertical aggregate supply curve satisfies the classical dichotomy, because it implies that the level of output is independent of the money supply. This long-run level of output, $\bar{Y}$, is called the *full-employment* or *natural* level of output. It is the level of output at which the economy's resources are fully employed or, more realistically, at which unemployment is at its natural rate.

figure 9-5



Shifts in Aggregate Demand in the Long Run


A reduction in the money supply shifts the aggregate demand curve downward from AD_1 to AD_2 . The equilibrium for the economy moves from point A to point B. Since the aggregate supply curve is vertical in the long run, the reduction in aggregate demand affects the price level but not the level of output.

The Short Run: The Horizontal Aggregate Supply Curve

The classical model and the vertical aggregate supply curve apply only in the long run. In the short run, some prices are sticky and, therefore, do not adjust to changes in demand. Because of this price stickiness, the short-run aggregate supply curve is not vertical.

As an extreme example, suppose that all firms have issued price catalogs and that it is costly for them to issue new ones. Thus, all prices are stuck at predetermined levels. At these prices, firms are willing to sell as much as their customers are willing to buy, and they hire just enough labor to produce the amount demanded. Because the price level is fixed, we represent this situation in Figure 9-6 with a horizontal aggregate supply curve.

figure 9-6



The Short-Run Aggregate Supply Curve In this extreme example, all prices are fixed in the short run. Therefore, the short-run aggregate supply curve, SRAS, is horizontal.


The short-run equilibrium of the economy is the intersection of the aggregate demand curve and this horizontal short-run aggregate supply curve. In this case, changes in aggregate demand do affect the level of output. For example, if the Fed suddenly reduces the money supply, the aggregate demand curve shifts inward, as in Figure 9-7. The economy moves from the old intersection of aggregate demand and aggregate supply, point A, to the new intersection, point B. The movement from point A to point B represents a decline in output at a fixed price level.

Thus, a fall in aggregate demand reduces output in the short run because prices do not adjust instantly. After the sudden fall in aggregate demand, firms are stuck with prices that are too high. With demand low and prices high, firms sell less of their product, so they reduce production and lay off workers. The economy experiences a recession.

From the Short Run to the Long Run

We can summarize our analysis so far as follows: *Over long periods of time, prices are flexible, the aggregate supply curve is vertical, and changes in aggregate demand affect the price level but not output. Over short periods of time, prices are sticky, the aggregate supply curve is flat, and changes in aggregate demand do affect the economy's output of goods and services.*


figure 9-7



Shifts in Aggregate Demand in the Short Run A reduction in the money supply shifts the aggregate demand curve downward from AD_1 to AD_2 . The equilibrium for the economy moves from point A to point B. Since the aggregate supply curve is horizontal in the short run, the reduction in aggregate demand reduces the level of output.


How does the economy make the transition from the short run to the long run? Let's trace the effects over time of a fall in aggregate demand. Suppose that the economy is initially in long-run equilibrium, as shown in Figure 9-8. In this figure, there are three curves: the aggregate demand curve, the long-run aggregate supply curve, and the short-run aggregate supply curve. The long-run equilibrium is the point at which aggregate demand crosses the long-run aggregate supply curve. Prices have adjusted to reach this equilibrium. Therefore, when the

figure 9-8



Long-Run Equilibrium In the long run, the economy finds itself at the intersection of the long-run aggregate supply curve and the aggregate demand curve. Because prices have adjusted to this level, the short-run aggregate supply curve crosses this point as well.

figure 9-9

**A Reduction in Aggregate Demand**

The economy begins in long-run equilibrium at point A. A reduction in aggregate demand, perhaps caused by a decrease in the money supply, moves the economy from point A to point B, where output is below its natural level. As prices fall, the economy gradually recovers from the recession, moving from point B to point C.

economy is in its long-run equilibrium, the short-run aggregate supply curve must cross this point as well.

Now suppose that the Fed reduces the money supply and the aggregate demand curve shifts downward, as in Figure 9-9. In the short run, prices are sticky, so the economy moves from point A to point B. Output and employment fall below their natural levels, which means the economy is in a recession. Over time, in response to the low demand, wages and prices fall. The gradual reduction in the price level moves the economy downward along the aggregate demand curve to point C, which is the new long-run equilibrium. In the new long-run equilibrium (point C), output and employment are back to their natural levels, but prices are lower than in the old long-run equilibrium (point A). Thus, a shift in aggregate demand affects output in the short run, but this effect dissipates over time as firms adjust their prices.

CASE STUDY**Gold, Greenbacks, and the Contraction of the 1870s**

The aftermath of the Civil War in the United States provides a vivid example of how contractionary monetary policy affects the economy. Before the war, the United States was on a gold standard. Paper dollars were readily convertible into gold. Under this policy, the quantity of gold determined the money supply and the price level.

In 1862, after the Civil War broke out, the Treasury announced that it would no longer redeem dollars for gold. In essence, this act replaced the gold standard with a system of fiat money. Over the next few years, the government printed large quantities of paper currency—called *greenbacks* for their color—and used

the seigniorage to finance wartime expenditure. Because of this increase in the money supply, the price level approximately doubled during the war.

When the war was over, much political debate centered on the question of whether to return to the gold standard. The Greenback party was formed with the primary goal of maintaining the system of fiat money. Eventually, however, the Greenback party lost the debate. Policymakers decided to retire the greenbacks over time in order to reinstate the gold standard at the rate of exchange between dollars and gold that had prevailed before the war. Their goal was to return the value of the dollar to its former level.

Returning to the gold standard in this way required reversing the wartime rise in prices, which meant aggregate demand had to fall. (To be more precise, the growth in aggregate demand needed to fall short of the growth in the natural rate of output.) As the price level fell, the economy experienced a recession from 1873 to 1879, the longest on record. By 1879, the price level was back to its level before the war, and the gold standard was reinstated.

9-4 Stabilization Policy

Fluctuations in the economy as a whole come from changes in aggregate supply or aggregate demand. Economists call exogenous changes in these curves **shocks** to the economy. A shock that shifts the aggregate demand curve is called a **demand shock**, and a shock that shifts the aggregate supply curve is called a **supply shock**. These shocks disrupt economic well-being by pushing output and employment away from their natural rates. One goal of the model of aggregate supply and aggregate demand is to show how shocks cause economic fluctuations.

Another goal of the model is to evaluate how macroeconomic policy can respond to these shocks. Economists use the term **stabilization policy** to refer to policy actions aimed at reducing the severity of short-run economic fluctuations. Because output and employment fluctuate around their long-run natural rates, stabilization policy dampens the business cycle by keeping output and employment as close to their natural rates as possible.

In the coming chapters, we examine in detail how stabilization policy works and what practical problems arise in its use. Here we begin our analysis of stabilization policy by examining how monetary policy might respond to shocks. Monetary policy is an important component of stabilization policy because, as we have seen, the money supply has a powerful impact on aggregate demand.


Shocks to Aggregate Demand

Consider an example of a demand shock: the introduction and expanded availability of credit cards. Because credit cards are often a more convenient way to make purchases than using cash, they reduce the quantity of money that people choose to hold. This reduction in money demand is equivalent to an increase in

the velocity of money. When each person holds less money, the money demand parameter k falls. This means that each dollar of money moves from hand to hand more quickly, so velocity $V (= 1/k)$ rises.

If the money supply is held constant, the increase in velocity causes nominal spending to rise and the aggregate demand curve to shift outward, as in Figure 9-10. In the short run, the increase in demand raises the output of the economy—it causes an economic boom. At the old prices, firms now sell more output. Therefore, they hire more workers, ask their existing workers to work longer hours, and make greater use of their factories and equipment.

figure 9-10



An Increase in Aggregate Demand

The economy begins in long-run equilibrium at point A. An increase in aggregate demand, due to an increase in the velocity of money, moves the economy from point A to point B, where output is above its natural level. As prices rise, output gradually returns to its natural rate, and the economy moves from point B to point C.

Over time, the high level of aggregate demand pulls up wages and prices. As the price level rises, the quantity of output demanded declines, and the economy gradually approaches the natural rate of production. But during the transition to the higher price level, the economy's output is higher than the natural rate.

What can the Fed do to dampen this boom and keep output closer to the natural rate? The Fed might reduce the money supply to offset the increase in velocity. Offsetting the change in velocity would stabilize aggregate demand. Thus, the Fed can reduce or even eliminate the impact of demand shocks on output and employment if it can skillfully control the money supply. Whether the Fed in fact has the necessary skill is a more difficult question, which we take up in Chapter 14.

Shocks to Aggregate Supply

Shocks to aggregate supply, as well as shocks to aggregate demand, can cause economic fluctuations. A supply shock is a shock to the economy that alters the cost of producing goods and services and, as a result, the prices that firms charge.

Because supply shocks have a direct impact on the price level, they are sometimes called *price shocks*. Here are some examples:


- A drought that destroys crops. The reduction in food supply pushes up food prices.
- A new environmental protection law that requires firms to reduce their emissions of pollutants. Firms pass on the added costs to customers in the form of higher prices.
- An increase in union aggressiveness. This pushes up wages and the prices of the goods produced by union workers.
- The organization of an international oil cartel. By curtailing competition, the major oil producers can raise the world price of oil.

All these events are *adverse* supply shocks, which means they push costs and prices upward. A *favorable* supply shock, such as the breakup of an international oil cartel, reduces costs and prices.

Figure 9-11 shows how an adverse supply shock affects the economy. The short-run aggregate supply curve shifts upward. (The supply shock may also lower the natural level of output and thus shift the long-run aggregate supply curve to the left, but we ignore that effect here.) If aggregate demand is held constant, the economy moves from point A to point B: the price level rises and the amount of output falls below the natural rate. An experience like this is called *stagflation*, because it combines stagnation (falling output) with inflation (rising prices).


Faced with an adverse supply shock, a policymaker controlling aggregate demand, such as the Fed, has a difficult choice between two options. The first option, implicit in Figure 9-11, is to hold aggregate demand constant. In this case, output and employment are lower than the natural rate. Eventually, prices

figure 9-11



An Adverse Supply Shock An adverse supply shock pushes up costs and thus prices. If aggregate demand is held constant, the economy moves from point A to point B, leading to stagflation—a combination of increasing prices and falling output. Eventually, as prices fall, the economy returns to the natural rate, point A.

figure 9-12



Accommodating an Adverse Supply Shock In response to an adverse supply shock, the Fed can increase aggregate demand to prevent a reduction in output. The economy moves from point A to point C. The cost of this policy is a permanently higher level of prices.

will fall to restore full employment at the old price level (point A). But the cost of this adjustment process is a painful recession.

The second option, illustrated in Figure 9-12, is to expand aggregate demand to bring the economy toward the natural rate more quickly. If the increase in aggregate demand coincides with the shock to aggregate supply, the economy goes immediately from point A to point C. In this case, the Fed is said to *accommodate* the supply shock. The drawback of this option, of course, is that the price level is permanently higher. There is no way to adjust aggregate demand to maintain full employment and keep the price level stable.

CASE STUDY

How OPEC Helped Cause Stagflation in the 1970s and Euphoria in the 1980s

The most disruptive supply shocks in recent history were caused by OPEC, the Organization of Petroleum Exporting Countries. In the early 1970s, OPEC's coordinated reduction in the supply of oil nearly doubled the world price. This increase in oil prices caused stagflation in most industrial countries. These statistics show what happened in the United States:

Year	Change in Oil Prices	Inflation Rate (CPI)	Unemployment Rate
1973	11.0%	6.2%	4.9%
1974	68.0	11.0	5.6
1975	16.0	9.1	8.5
1976	3.3	5.8	7.7
1977	8.1	6.5	7.1

The 68-percent increase in the price of oil in 1974 was an adverse supply shock of major proportions. As one would have expected, it led to both higher inflation and higher unemployment.

A few years later, when the world economy had nearly recovered from the first OPEC recession, almost the same thing happened again. OPEC raised oil prices, causing further stagflation. Here are the statistics for the United States:

Year	Change in Oil Prices	Inflation Rate (CPI)	Unemployment Rate
1978	9.4%	7.7%	6.1%
1979	25.4	11.3	5.8
1980	47.8	13.5	7.0
1981	44.4	10.3	7.5
1982	−8.7	6.1	9.5

The increases in oil prices in 1979, 1980, and 1981 again led to double-digit inflation and higher unemployment.

In the mid-1980s, political turmoil among the Arab countries weakened OPEC's ability to restrain supplies of oil. Oil prices fell, reversing the stagflation of the 1970s and the early 1980s. Here's what happened:

Year	Change in Oil Prices	Inflation Rate (CPI)	Unemployment Rate
1983	−7.1%	3.2%	9.5%
1984	−1.7	4.3	7.4
1985	−7.5	3.6	7.1
1986	−44.5	1.9	6.9
1987	18.3	3.6	6.1

In 1986 oil prices fell by nearly half. This favorable supply shock led to one of the lowest inflation rates experienced in recent U.S. history and to falling unemployment.

More recently, OPEC has not been a major cause of economic fluctuations. This is in part because OPEC has been less successful at raising the price of oil. Although world oil prices have fluctuated, the changes have not been as large as those experienced during the 1970s, and the real price of oil has never returned to the peaks reached in the early 1980s. Moreover, conservation efforts and technological changes have made the economy less susceptible to oil shocks. The amount of oil consumed per unit of real GDP has fallen about 40 percent over the past three decades.

But we should not be too sanguine. The experiences of the 1970s and 1980s could always be repeated. Events in the Middle East are a potential source of shocks to economies around the world.³

³ Some economists have suggested that changes in oil prices played a major role in economic fluctuations even before the 1970s. See James D. Hamilton, "Oil and the Macroeconomy Since World War II," *Journal of Political Economy* 91 (April 1983): 228–248.

9-5 Conclusion

This chapter introduced a framework to study economic fluctuations: the model of aggregate supply and aggregate demand. The model is built on the assumption that prices are sticky in the short run and flexible in the long run. It shows how shocks to the economy cause output to deviate temporarily from the level implied by the classical model.

The model also highlights the role of monetary policy. Poor monetary policy can be a source of shocks to the economy. A well-run monetary policy can respond to shocks and stabilize the economy.

In the chapters that follow, we refine our understanding of this model and our analysis of stabilization policy. Chapters 10 through 12 go beyond the quantity equation to refine our theory of aggregate demand. This refinement shows that aggregate demand depends on fiscal policy as well as monetary policy. Chapter 13 examines aggregate supply in more detail. Chapter 14 examines the debate over the virtues and limits of stabilization policy.

Summary

1. The crucial difference between the long run and the short run is that prices are flexible in the long run but sticky in the short run. The model of aggregate supply and aggregate demand provides a framework to analyze economic fluctuations and see how the impact of policies varies over different time horizons.
2. The aggregate demand curve slopes downward. It tells us that the lower the price level, the greater the aggregate quantity of goods and services demanded.
3. In the long run, the aggregate supply curve is vertical because output is determined by the amounts of capital and labor and by the available technology, but not by the level of prices. Therefore, shifts in aggregate demand affect the price level but not output or employment.
4. In the short run, the aggregate supply curve is horizontal, because wages and prices are sticky at predetermined levels. Therefore, shifts in aggregate demand affect output and employment.
5. Shocks to aggregate demand and aggregate supply cause economic fluctuations. Because the Fed can shift the aggregate demand curve, it can attempt to offset these shocks to maintain output and employment at their natural rates.

KEY CONCEPTS

Aggregate demand
Aggregate supply

Shocks
Demand shocks

Supply shocks
Stabilization policy

QUESTIONS FOR REVIEW

1. Give an example of a price that is sticky in the short run and flexible in the long run.
2. Why does the aggregate demand curve slope downward?
3. Explain the impact of an increase in the money supply in the short run and in the long run.
4. Why is it easier for the Fed to deal with demand shocks than with supply shocks?

PROBLEMS AND APPLICATIONS

1. Suppose that a change in government regulations allows banks to start paying interest on checking accounts. Recall that the money stock is the sum of currency and demand deposits, including checking accounts, so this regulatory change makes holding money more attractive.
 - a. How does this change affect the demand for money?
 - b. What happens to the velocity of money?
 - c. If the Fed keeps the money supply constant, what will happen to output and prices in the short run and in the long run?
 - d. Should the Fed keep the money supply constant in response to this regulatory change? Why or why not?
2. Suppose the Fed reduces the money supply by 5 percent.
 - a. What happens to the aggregate demand curve?
 - b. What happens to the level of output and the price level in the short run and in the long run?
 - c. According to Okun's law, what happens to unemployment in the short run and in the long run? (*Hint:* Okun's law is the relationship between output and unemployment discussed in Chapter 2.)
 - d. What happens to the real interest rate in the short run and in the long run? (*Hint:* Use the model of the real interest rate in Chapter 3 to see what happens when output changes.)
3. Let's examine how the goals of the Fed influence its response to shocks. Suppose Fed A cares only about keeping the price level stable, and Fed B cares only about keeping output and employment at their natural rates. Explain how each Fed would respond to
 - a. An exogenous decrease in the velocity of money.
 - b. An exogenous increase in the price of oil.
4. The official arbiter of when recessions begin and end is the National Bureau of Economic Research, a nonprofit economics research group. Go to the NBER's Web site (www.nber.org) and find the latest turning point in the business cycle. When did it occur? Was this a switch from expansion to contraction or the other way around? List all the recessions (contractions) that have occurred during your lifetime and the dates when they began and ended.

Aggregate Demand I

I shall argue that the postulates of the classical theory are applicable to a special case only and not to the general case. . . . Moreover, the characteristics of the special case assumed by the classical theory happen not to be those of the economic society in which we actually live, with the result that its teaching is misleading and disastrous if we attempt to apply it to the facts of experience.

— John Maynard Keynes, *The General Theory*

Of all the economic fluctuations in world history, the one that stands out as particularly large, painful, and intellectually significant is the Great Depression of the 1930s. During this time, the United States and many other countries experienced massive unemployment and greatly reduced incomes. In the worst year, 1933, one-fourth of the U.S. labor force was unemployed, and real GDP was 30 percent below its 1929 level.

This devastating episode caused many economists to question the validity of classical economic theory—the theory we examined in Chapters 3 through 6. Classical theory seemed incapable of explaining the Depression. According to that theory, national income depends on factor supplies and the available technology, neither of which changed substantially from 1929 to 1933. After the onset of the Depression, many economists believed that a new model was needed to explain such a large and sudden economic downturn and to suggest government policies that might reduce the economic hardship so many people faced.

In 1936 the British economist John Maynard Keynes revolutionized economics with his book *The General Theory of Employment, Interest, and Money*. Keynes proposed a new way to analyze the economy, which he presented as an alternative to classical theory. His vision of how the economy works quickly became a center of controversy. Yet, as economists debated *The General Theory*, a new understanding of economic fluctuations gradually developed.

Keynes proposed that low aggregate demand is responsible for the low income and high unemployment that characterize economic downturns. He criticized classical theory for assuming that aggregate supply alone—capital, labor, and technology—determines national income. Economists today reconcile these two views


with the model of aggregate demand and aggregate supply introduced in Chapter 9. In the long run, prices are flexible, and aggregate supply determines income. But in the short run, prices are sticky, so changes in aggregate demand influence income.

In this chapter and the next, we continue our study of economic fluctuations by looking more closely at aggregate demand. Our goal is to identify the variables that shift the aggregate demand curve, causing fluctuations in national income. We also examine more fully the tools policymakers can use to influence aggregate demand. In Chapter 9 we derived the aggregate demand curve from the quantity theory of money, and we showed that monetary policy can shift the aggregate demand curve. In this chapter we see that the government can influence aggregate demand with both monetary and fiscal policy.

The model of aggregate demand developed in this chapter, called the **IS–LM model**, is the leading interpretation of Keynes’s theory. The goal of the model is to show what determines national income for any given price level. There are two ways to view this exercise. We can view the *IS–LM* model as showing what causes income to change in the short run when the price level is fixed. Or we can view the model as showing what causes the aggregate demand curve to shift. These two views of the model are equivalent: as Figure 10-1 shows, in the short run when the price level is fixed, shifts in the aggregate demand curve lead to changes in national income.

The two parts of the *IS–LM* model are, not surprisingly, the **IS curve** and the **LM curve**. *IS* stands for “investment” and “saving,” and the *IS* curve represents what’s going on in the market for goods and services (which we first discussed in Chapter 3). *LM* stands for “liquidity” and “money,” and the *LM* curve represents what’s happening to the supply and demand for money (which we first discussed in Chapter 4). Because the interest rate influences both investment and money

figure 10-1



Shifts in Aggregate Demand For a given price level, national income fluctuates because of shifts in the aggregate demand curve. The *IS–LM* model takes the price level as given and shows what causes income to change. The model therefore shows what causes aggregate demand to shift.

demand, it is the variable that links the two halves of the *IS–LM* model. The model shows how interactions between these markets determine the position and slope of the aggregate demand curve and, therefore, the level of national income in the short run.¹

10-1 The Goods Market and the *IS* Curve

The *IS* curve plots the relationship between the interest rate and the level of income that arises in the market for goods and services. To develop this relationship, we start with a basic model called the **Keynesian cross**. This model is the simplest interpretation of Keynes's theory of national income and is a building block for the more complex and realistic *IS–LM* model.

The Keynesian Cross

In *The General Theory*, Keynes proposed that an economy's total income was, in the short run, determined largely by the desire to spend by households, firms, and the government. The more people want to spend, the more goods and services firms can sell. The more firms can sell, the more output they will choose to produce and the more workers they will choose to hire. Thus, the problem during recessions and depressions, according to Keynes, was inadequate spending. The Keynesian cross is an attempt to model this insight.

Planned Expenditure We begin our derivation of the Keynesian cross by drawing a distinction between actual and planned expenditure. *Actual expenditure* is the amount households, firms, and the government spend on goods and services, and as we first saw in Chapter 2, it equals the economy's gross domestic product (GDP). *Planned expenditure* is the amount households, firms, and the government would like to spend on goods and services.

Why would actual expenditure ever differ from planned expenditure? The answer is that firms might engage in unplanned inventory investment because their sales do not meet their expectations. When firms sell less of their product than they planned, their stock of inventories automatically rises; conversely, when firms sell more than planned, their stock of inventories falls. Because these unplanned changes in inventory are counted as investment spending by firms, actual expenditure can be either above or below planned expenditure.

Now consider the determinants of planned expenditure. Assuming that the economy is closed, so that net exports are zero, we write planned expenditure *E* as the sum of consumption *C*, planned investment *I*, and government purchases *G*:

$$E = C + I + G.$$

¹ The *IS–LM* model was introduced in a classic article by the Nobel-Prize-winning economist John R. Hicks, "Mr. Keynes and the Classics: A Suggested Interpretation," *Econometrica* 5 (1937): 147–159.

To this equation, we add the consumption function

$$C = C(Y - T).$$

This equation states that consumption depends on disposable income ($Y - T$), which is total income Y minus taxes T . To keep things simple, for now we take planned investment as exogenously fixed:

$$I = \bar{I}.$$

And as in Chapter 3, we assume that fiscal policy—the levels of government purchases and taxes—is fixed:

$$G = \bar{G},$$

$$T = \bar{T}.$$

Combining these five equations, we obtain

$$E = C(Y - \bar{T}) + \bar{I} + \bar{G}.$$


This equation shows that planned expenditure is a function of income Y , the level of planned investment $\bar{I}$, and the fiscal policy variables $\bar{G}$ and $\bar{T}$.

Figure 10-2 graphs planned expenditure as a function of the level of income. This line slopes upward because higher income leads to higher consumption and thus higher planned expenditure. The slope of this line is the marginal propensity to consume, the *MPC*: it shows how much planned expenditure increases when income rises by \$1. This planned-expenditure function is the first piece of the model called the Keynesian cross.

The Economy in Equilibrium The next piece of the Keynesian cross is the assumption that the economy is in equilibrium when actual expenditure equals planned expenditure. This assumption is based on the idea that when people's plans have been realized, they have no reason to change what they are doing.

figure 10-2

Planned expenditure, E



Planned Expenditure as a Function of Income Planned expenditure depends on income because higher income leads to higher consumption, which is part of planned expenditure. The slope of this planned-expenditure function is the marginal propensity to consume, *MPC*.


Recalling that Y as GDP equals not only total income but also total actual expenditure on goods and services, we can write this equilibrium condition as

$$\begin{aligned}\text{Actual Expenditure} &= \text{Planned Expenditure} \\ Y &= E.\end{aligned}$$

The 45-degree line in Figure 10-3 plots the points where this condition holds. With the addition of the planned-expenditure function, this diagram becomes the Keynesian cross. The equilibrium of this economy is at point A, where the planned-expenditure function crosses the 45-degree line.

How does the economy get to the equilibrium? In this model, inventories play an important role in the adjustment process. Whenever the economy is not in equilibrium, firms experience unplanned changes in inventories, and this induces them to change production levels. Changes in production in turn influence total income and expenditure, moving the economy toward equilibrium.

figure 10-3




The Keynesian Cross The equilibrium in the Keynesian cross is at point A, where income (actual expenditure) equals planned expenditure.

For example, suppose the economy were ever to find itself with GDP at a level greater than the equilibrium level, such as the level Y_1 in Figure 10-4. In this case, planned expenditure E_1 is less than production Y_1 , so firms are selling less than they are producing. Firms add the unsold goods to their stock of inventories. This unplanned rise in inventories induces firms to lay off workers and reduce production, and these actions in turn reduce GDP. This process of unintended inventory accumulation and falling income continues until income Y falls to the equilibrium level.

Similarly, suppose GDP were at a level lower than the equilibrium level, such as the level Y_2 in Figure 10-4. In this case, planned expenditure E_2 is greater than production Y_2 . Firms meet the high level of sales by drawing down their inventories.

figure 10-4



The Adjustment to Equilibrium in the Keynesian Cross If firms were producing at level Y_1 , then planned expenditure E_1 would fall short of production, and firms would accumulate inventories. This inventory accumulation would induce firms to reduce production. Similarly, if firms were producing at level Y_2 , then planned expenditure E_2 would exceed production, and firms would run down their inventories. This fall in inventories would induce firms to raise production. In both cases, the firms' decisions drive the economy toward equilibrium.

But when firms see their stock of inventories dwindle, they hire more workers and increase production. GDP rises, and the economy approaches the equilibrium.

In summary, the Keynesian cross shows how income Y is determined for given levels of planned investment I and fiscal policy G and T . We can use this model to show how income changes when one of these exogenous variables changes.

Fiscal Policy and the Multiplier: Government Purchases Consider how changes in government purchases affect the economy. Because government purchases are one component of expenditure, higher government purchases result in higher planned expenditure for any given level of income. If government purchases rise by ΔG , then the planned-expenditure schedule shifts upward by ΔG , as in Figure 10-5. The equilibrium of the economy moves from point A to point B.

This graph shows that an increase in government purchases leads to an even greater increase in income. That is, ΔY is larger than ΔG . The ratio $\Delta Y / \Delta G$ is called the **government-purchases multiplier**; it tells us how much income rises in response to a \$1 increase in government purchases. An implication of the Keynesian cross is that the government-purchases multiplier is larger than 1.


Why does fiscal policy have a multiplied effect on income?



© The New Yorker Collection, 1992 Dana Fradon from cartoonbank.com. All Rights Reserved.

"Your Majesty, my voyage will not only forge a new route to the spices of the East but also create over three thousand new jobs."

figure 10-5



An Increase in Government Purchases in the Keynesian Cross

An increase in government purchases of ΔG raises planned expenditure by that amount for any given level of income. The equilibrium moves from point A to point B, and income rises from Y_1 to Y_2 . Note that the increase in income ΔY exceeds the increase in government purchases ΔG . Thus, fiscal policy has a multiplied effect on income.

The reason is that, according to the consumption function $C = C(Y - T)$, higher income causes higher consumption. When an increase in government purchases raises income, it also raises consumption, which further raises income, which further raises consumption, and so on. Therefore, in this model, an increase in government purchases causes a greater increase in income.

How big is the multiplier? To answer this question, we trace through each step of the change in income. The process begins when expenditure rises by ΔG , which implies that income rises by ΔG as well. This increase in income in turn raises consumption by $MPC \times \Delta G$, where MPC is the marginal propensity to consume. This increase in consumption raises expenditure and income once again. This second increase in income of $MPC \times \Delta G$ again raises consumption, this time by $MPC \times (MPC \times \Delta G)$, which again raises expenditure and income, and so on. This feedback from consumption to income to consumption continues indefinitely. The total effect on income is

Initial Change in Government Purchases =	ΔG
First Change in Consumption	$= MPC \times \Delta G$
Second Change in Consumption	$= MPC^2 \times \Delta G$
Third Change in Consumption	$= MPC^3 \times \Delta G$
⋮	⋮
⋮	⋮
$\Delta Y = (1 + MPC + MPC^2 + MPC^3 + \dots) \Delta G.$	

The government-purchases multiplier is

$$\Delta Y / \Delta G = 1 + MPC + MPC^2 + MPC^3 + \dots$$

This expression for the multiplier is an example of an *infinite geometric series*. A result from algebra allows us to write the multiplier as²

$$\Delta Y/\Delta G = 1/(1 - MPC).$$

For example, if the marginal propensity to consume is 0.6, the multiplier is

$$\begin{aligned}\Delta Y/\Delta G &= 1 + 0.6 + 0.6^2 + 0.6^3 + \dots \\ &= 1/(1 - 0.6) \\ &= 2.5.\end{aligned}$$

In this case, a \$1.00 increase in government purchases raises equilibrium income by \$2.50.³

Fiscal Policy and the Multiplier: Taxes Consider now how changes in taxes affect equilibrium income. A decrease in taxes of ΔT immediately raises disposable income $Y - T$ by ΔT and, therefore, increases consumption by $MPC \times \Delta T$. For any given level of income Y , planned expenditure is now higher. As Figure 10-6 shows, the planned-expenditure schedule shifts upward by $MPC \times \Delta T$. The equilibrium of the economy moves from point A to point B.

² *Mathematical note:* We prove this algebraic result as follows. Let

$$z = 1 + x + x^2 + \dots,$$

Multiply both sides of this equation by x :

$$xz = x + x^2 + x^3 + \dots,$$

Subtract the second equation from the first:

$$z - xz = 1.$$

Rearrange this last equation to obtain

$$z(1 - x) = 1,$$

which implies

$$z = 1/(1 - x).$$

This completes the proof.

³ *Mathematical note:* The government-purchases multiplier is most easily derived using a little calculus. Begin with the equation

$$Y = C(Y - T) + I + G.$$

Holding T and I fixed, differentiate to obtain


$$dY = C' dY + dG,$$

and then rearrange to find

$$dY/dG = 1/(1 - C').$$

This is the same as the equation in the text.

figure 10-6



A Decrease in Taxes in the Keynesian Cross A decrease in taxes of ΔT raises planned expenditure by $MPC \times \Delta T$ for any given level of income. The equilibrium moves from point A to point B, and income rises from Y_1 to Y_2 . Again, fiscal policy has a multiplied effect on income.

Just as an increase in government purchases has a multiplied effect on income, so does a decrease in taxes. As before, the initial change in expenditure, now $MPC \times \Delta T$, is multiplied by $1/(1 - MPC)$. The overall effect on income of the change in taxes is

$$\Delta Y / \Delta T = -MPC / (1 - MPC).$$

This expression is the **tax multiplier**, the amount income changes in response to a \$1 change in taxes. For example, if the marginal propensity to consume is 0.6, then the tax multiplier is

$$\Delta Y / \Delta T = -0.6 / (1 - 0.6) = -1.5.$$

In this example, a \$1.00 cut in taxes raises equilibrium income by \$1.50.⁴

⁴ *Mathematical note:* As before, the multiplier is most easily derived using a little calculus. Begin with the equation

$$Y = C(Y - T) + I + G.$$

Holding I and G fixed, differentiate to obtain

$$dY = C'(dY - dT),$$

and then rearrange to find

$$dY/dT = -C'/(1 - C').$$

This is the same as the equation in the text.

CASE STUDY

Cutting Taxes to Stimulate the Economy

When John F. Kennedy became president of the United States in 1961, he brought to Washington some of the brightest young economists of the day to work on his Council of Economic Advisers. These economists, who had been schooled in the economics of Keynes, brought Keynesian ideas to discussions of economic policy at the highest level.

One of the council's first proposals was to expand national income by reducing taxes. This eventually led to a substantial cut in personal and corporate income taxes in 1964. The tax cut was intended to stimulate expenditure on consumption and investment and thus lead to higher levels of income and employment. When a reporter asked Kennedy why he advocated a tax cut, Kennedy replied, "To stimulate the economy. Don't you remember your Economics 101?"

As Kennedy's economic advisers predicted, the passage of the tax cut was followed by an economic boom. Growth in real GDP was 5.3 percent in 1964 and 6.0 percent in 1965. The unemployment rate fell from 5.7 percent in 1963 to 5.2 percent in 1964 and then to 4.5 percent in 1965.⁵

Economists continue to debate the source of this rapid growth in the early 1960s. A group called *supply-siders* argues that the economic boom resulted from the incentive effects of the cut in income tax rates. According to supply-siders, when workers are allowed to keep a higher fraction of their earnings, they supply substantially more labor and expand the aggregate supply of goods and services. Keynesians, however, emphasize the impact of tax cuts on aggregate demand. Most likely, both views have some truth: *Tax cuts stimulate aggregate supply by improving workers' incentives and expand aggregate demand by raising households' disposable income.*

When George W. Bush was elected president in 2001, a major element of his platform was a cut in income taxes. Bush and his advisers used both supply-side and Keynesian rhetoric to make the case for their policy. During the campaign, when the economy was doing fine, they argued that lower marginal tax rates would improve work incentives. But then the economy started to slow: unemployment rose from 3.9 percent in October 2000 to 4.5 percent in April 2001. The argument shifted to emphasize that the tax cut would stimulate spending and reduce the risk of recession.

Congress passed the tax cut in May 2001. Compared to the original Bush proposal, the bill cut tax rates less in the long run. But it added an immediate tax rebate of \$600 per family (\$300 for single taxpayers) that was mailed out in the summer of 2001. Consistent with Keynesian theory, the goal of the rebate was to provide an immediate stimulus to aggregate demand.

⁵ For an analysis of the 1964 tax cut by one of Kennedy's economists, see Arthur Okun, "Measuring the Impact of the 1964 Tax Reduction," in W. W. Heller, ed., *Perspectives on Economic Growth* (New York: Random House, 1968); reprinted in Arthur M. Okun, *Economics for Policymaking* (Cambridge, MA: MIT Press, 1983), 405–423.

The Interest Rate, Investment, and the *IS* Curve

The Keynesian cross is only a steppingstone on our path to the *IS*–*LM* model. The Keynesian cross is useful because it shows how the spending plans of households, firms, and the government determine the economy's income. Yet it makes the simplifying assumption that the level of planned investment I is fixed. As we discussed in Chapter 3, an important macroeconomic relationship is that planned investment depends on the interest rate r .


To add this relationship between the interest rate and investment to our model, we write the level of planned investment as

$$I = I(r).$$

This investment function is graphed in panel (a) of Figure 10-7. Because the interest rate is the cost of borrowing to finance investment projects, an increase in

figure 10-7

Deriving the *IS* Curve Panel (a) shows the investment function: an increase in the interest rate from r_1 to r_2 reduces planned investment from $I(r_1)$ to $I(r_2)$. Panel (b) shows the Keynesian cross: a decrease in planned investment from $I(r_1)$ to $I(r_2)$ shifts the planned-expenditure function downward and thereby reduces income from Y_1 to Y_2 . Panel (c) shows the *IS* curve summarizing this relationship between the interest rate and income: the higher the interest rate, the lower the level of income.



the interest rate reduces planned investment. As a result, the investment function slopes downward.

To determine how income changes when the interest rate changes, we can combine the investment function with the Keynesian-cross diagram. Because investment is inversely related to the interest rate, an increase in the interest rate from r_1 to r_2 reduces the quantity of investment from $I(r_1)$ to $I(r_2)$. The reduction in planned investment, in turn, shifts the planned-expenditure function downward, as in panel (b) of Figure 10-7. The shift in the planned-expenditure function causes the level of income to fall from Y_1 to Y_2 . Hence, an increase in the interest rate lowers income.

The *IS* curve, shown in panel (c) of Figure 10-7, summarizes this relationship between the interest rate and the level of income. In essence, the *IS* curve combines the interaction between r and I expressed by the investment function and the interaction between I and Y demonstrated by the Keynesian cross. Because an increase in the interest rate causes planned investment to fall, which in turn causes income to fall, the *IS* curve slopes downward.

How Fiscal Policy Shifts the *IS* Curve

The *IS* curve shows us, for any given interest rate, the level of income that brings the goods market into equilibrium. As we learned from the Keynesian cross, the level of income also depends on fiscal policy. The *IS* curve is drawn for a given fiscal policy; that is, when we construct the *IS* curve, we hold G and T fixed. When fiscal policy changes, the *IS* curve shifts.

Figure 10-8 uses the Keynesian cross to show how an increase in government purchases by ΔG shifts the *IS* curve. This figure is drawn for a given interest rate $\bar{r}$ and thus for a given level of planned investment. The Keynesian cross shows that this change in fiscal policy raises planned expenditure and thereby increases equilibrium income from Y_1 to Y_2 . Therefore, an increase in government purchases shifts the *IS* curve outward.


We can use the Keynesian cross to see how other changes in fiscal policy shift the *IS* curve. Because a decrease in taxes also expands expenditure and income, it too shifts the *IS* curve outward. A decrease in government purchases or an increase in taxes reduces income; therefore, such a change in fiscal policy shifts the *IS* curve inward.

In summary, the IS curve shows the combinations of the interest rate and the level of income that are consistent with equilibrium in the market for goods and services. The IS curve is drawn for a given fiscal policy. Changes in fiscal policy that raise the demand for goods and services shift the IS curve to the right. Changes in fiscal policy that reduce the demand for goods and services shift the IS curve to the left.

A Loanable-Funds Interpretation of the *IS* Curve

When we first studied the market for goods and services in Chapter 3, we noted an equivalence between the supply and demand for goods and services and the supply and demand for loanable funds. This equivalence provides another way to interpret the *IS* curve.

figure 10-8



An Increase in Government Purchases Shifts the IS Curve Outward

Panel (a) shows that an increase in government purchases raises planned expenditure. For any given interest rate, the upward shift in planned expenditure of ΔG leads to an increase in income Y of $\Delta G / (1 - MPC)$. Therefore, in panel (b), the IS curve shifts to the right by this amount.

Recall that the national income accounts identity can be written as

$$Y - C - G = I$$

$$S = I.$$

The left-hand side of this equation is national saving S , and the right-hand side is investment I . National saving represents the supply of loanable funds, and investment represents the demand for these funds.

To see how the market for loanable funds produces the *IS* curve, substitute the consumption function for C and the investment function for I :

$$Y - C(Y - T) - G = I(r).$$


The left-hand side of this equation shows that the supply of loanable funds depends on income and fiscal policy. The right-hand side shows that the demand for loanable funds depends on the interest rate. The interest rate adjusts to equilibrate the supply and demand for loans.

As Figure 10-9 illustrates, we can interpret the *IS* curve as showing the interest rate that equilibrates the market for loanable funds for any given level of income. When income rises from Y_1 to Y_2 , national saving, which equals $Y - C - G$, increases. (Consumption rises by less than income, because the marginal propensity to consume is less than 1.) As panel (a) shows, the increased supply of loanable funds drives down the interest rate from r_1 to r_2 . The *IS* curve in panel (b) summarizes this relationship: higher income implies higher saving, which in turn implies a lower equilibrium interest rate. For this reason, the *IS* curve slopes downward.

This alternative interpretation of the *IS* curve also explains why a change in fiscal policy shifts the *IS* curve. An increase in government purchases or a decrease in taxes reduces national saving for any given level of income. The reduced supply of loanable funds raises the interest rate that equilibrates the market. Because the interest rate is now higher for any given level of income, the *IS* curve shifts upward in response to the expansionary change in fiscal policy.

Finally, note that the *IS* curve does not determine either income Y or the interest rate r . Instead, the *IS* curve is a relationship between Y and r arising in the

figure 10-9



A Loanable-Funds Interpretation of the *IS* Curve Panel (a) shows that an increase in income from Y_1 to Y_2 raises saving and thus lowers the interest rate that equilibrates the supply and demand for loanable funds. The *IS* curve in panel (b) expresses this negative relationship between income and the interest rate.

market for goods and services or, equivalently, the market for loanable funds. To determine the equilibrium of the economy, we need another relationship between these two variables, to which we now turn.

10-2 The Money Market and the *LM* Curve

The *LM* curve plots the relationship between the interest rate and the level of income that arises in the market for money balances. To understand this relationship, we begin by looking at a theory of the interest rate, called the **theory of liquidity preference**.

The Theory of Liquidity Preference

In his classic work *The General Theory*, Keynes offered his view of how the interest rate is determined in the short run. That explanation is called the theory of liquidity preference, because it posits that the interest rate adjusts to balance the supply and demand for the economy's most liquid asset—money. Just as the Keynesian cross is a building block for the *IS* curve, the theory of liquidity preference is a building block for the *LM* curve.

To develop this theory, we begin with the supply of real money balances. If M stands for the supply of money and P stands for the price level, then M/P is the supply of real money balances. The theory of liquidity preference assumes there is a fixed supply of real money balances. That is,


$$(M/P)^s = \bar{M}/\bar{P}.$$

The money supply M is an exogenous policy variable chosen by a central bank, such as the Federal Reserve. The price level P is also an exogenous variable in this model. (We take the price level as given because the *IS–LM* model—our ultimate goal in this chapter—explains the short run when the price level is fixed.) These assumptions imply that the supply of real money balances is fixed and, in particular, does not depend on the interest rate. Thus, when we plot the supply of real money balances against the interest rate in Figure 10-10, we obtain a vertical supply curve.

Next, consider the demand for real money balances. The theory of liquidity preference posits that the interest rate is one determinant of how much money people choose to hold. The reason is that the interest rate is the opportunity cost of holding money: it is what you forgo by holding some of your assets as money, which does not bear interest, instead of as interest-bearing bank deposits or bonds. When the interest rate rises, people want to hold less of their wealth in the form of money. We can write the demand for real money balances as

$$(M/P)^d = L(r),$$

figure 10-10



The Theory of Liquidity Preference The supply and demand for real money balances determine the interest rate. The supply curve for real money balances is vertical because the supply does not depend on the interest rate. The demand curve is downward sloping because a higher interest rate raises the cost of holding money and thus lowers the quantity demanded. At the equilibrium interest rate, the quantity of real money balances demanded equals the quantity supplied.


where the function $L()$ shows that the quantity of money demanded depends on the interest rate. Thus, the demand curve in Figure 10-10 slopes downward because higher interest rates reduce the quantity of real money balances demanded.⁶

According to the theory of liquidity preference, the supply and demand for real money balances determine what interest rate prevails in the economy. That is, the interest rate adjusts to equilibrate the money market. As the figure shows, at the equilibrium interest rate, the quantity of real money balances demanded equals the quantity supplied.

How does the interest rate get to this equilibrium of money supply and money demand? The adjustment occurs because whenever the money market is not in equilibrium, people try to adjust their portfolios of assets and, in the process, alter the interest rate. For instance, if the interest rate is above the equilibrium level, the quantity of real money balances supplied exceeds the quantity demanded. Individuals holding the excess supply of money try to convert some of their non-interest-bearing money into interest-bearing bank deposits or bonds. Banks and bond issuers, who prefer to pay lower interest rates, respond to this excess supply of money by lowering the interest rates they offer. Conversely, if the interest rate is below the equilibrium level, so that the quantity of money demanded exceeds the quantity supplied, individuals try to obtain money by selling bonds or making bank withdrawals. To attract now-scarcer funds, banks and bond issuers respond by increasing the interest rates they offer. Eventually, the

⁶ Note that r is being used to denote the interest rate here, as it was in our discussion of the IS curve. More accurately, it is the nominal interest rate that determines money demand and the real interest rate that determines investment. To keep things simple, we are ignoring expected inflation, which creates the difference between the real and nominal interest rates. The role of expected inflation in the $IS-LM$ model is explored in Chapter 11.

figure 10-11



A Reduction in the Money Supply in the Theory of Liquidity Preference

If the price level is fixed, a reduction in the money supply from M_1 to M_2 reduces the supply of real money balances. The equilibrium interest rate therefore rises from r_1 to r_2 .

interest rate reaches the equilibrium level, at which people are content with their portfolios of monetary and nonmonetary assets.

Now that we have seen how the interest rate is determined, we can use the theory of liquidity preference to show how the interest rate responds to changes in the supply of money. Suppose, for instance, that the Fed suddenly decreases the money supply. A fall in M reduces M/P , because P is fixed in the model. The supply of real money balances shifts to the left, as in Figure 10-11. The equilibrium interest rate rises from r_1 to r_2 , and the higher interest rate makes people satisfied to hold the smaller quantity of real money balances. The opposite would occur if the Fed had suddenly increased the money supply. Thus, according to the theory of liquidity preference, a decrease in the money supply raises the interest rate, and an increase in the money supply lowers the interest rate.

CASE STUDY

Did Paul Volcker's Monetary Tightening Raise or Lower Interest Rates?

The early 1980s saw the largest and quickest reduction in inflation in recent U.S. history. By the late 1970s inflation had reached the double-digit range; in 1979, consumer prices were rising at a rate of 11.3 percent per year. In October 1979, only two months after becoming the chairman of the Federal Reserve, Paul Volcker announced that monetary policy would aim to reduce the rate of inflation. This announcement began a period of tight money that, by 1983, brought the inflation rate down to about 3 percent.

How does such a monetary tightening influence interest rates? According to the theories we have been developing, the answer depends on the time horizon. Our analysis of the Fisher effect in Chapter 4 suggests that in the long run

Volcker's change in monetary policy would lower inflation, and this in turn would lead to lower nominal interest rates. Yet the theory of liquidity preference predicts that, in the short run when prices are sticky, anti-inflationary monetary policy would lead to falling real money balances and higher nominal interest rates.

Both conclusions are consistent with experience. Nominal interest rates did fall in the 1980s as inflation fell. But comparing the year before the October 1979 announcement and the year after, we find that real money balances ($M1$ divided by the CPI) fell 8.3 percent and the nominal interest rate (on short-term commercial loans) rose from 10.1 percent to 11.9 percent. Hence, although a monetary tightening leads to lower nominal interest rates in the long run, it leads to higher nominal interest rates in the short run.

Income, Money Demand, and the LM Curve

Having developed the theory of liquidity preference as an explanation for what determines the interest rate, we can now use the theory to derive the LM curve. We begin by considering the following question: How does a change in the economy's level of income Y affect the market for real money balances? The answer (which should be familiar from Chapter 4) is that the level of income affects the demand for money. When income is high, expenditure is high, so people engage in more transactions that require the use of money. Thus, greater income implies greater money demand. We can express these ideas by writing the money demand function as

$$(M/P)^d = L(r, Y).$$

The quantity of real money balances demanded is negatively related to the interest rate and positively related to income.


Using the theory of liquidity preference, we can figure out what happens to the equilibrium interest rate when the level of income changes. For example, consider what happens in Figure 10-12 when income increases from Y_1 to Y_2 . As panel (a) illustrates, this increase in income shifts the money demand curve to the right. With the supply of real money balances unchanged, the interest rate must rise from r_1 to r_2 to equilibrate the money market. Therefore, according to the theory of liquidity preference, higher income leads to a higher interest rate.

The LM curve plots this relationship between the level of income and the interest rate. The higher the level of income, the higher the demand for real money balances, and the higher the equilibrium interest rate. For this reason, the LM curve slopes upward, as in panel (b) of Figure 10-12.

How Monetary Policy Shifts the LM Curve

The LM curve tells us the interest rate that equilibrates the money market at any level of income. Yet, as we saw earlier, the equilibrium interest rate also depends on the supply of real money balances, M/P . This means that the LM curve is

figure 10-12



Deriving the LM Curve Panel (a) shows the market for real money balances: an increase in income from Y_1 to Y_2 raises the demand for money and thus raises the interest rate from r_1 to r_2 . Panel (b) shows the LM curve summarizing this relationship between the interest rate and income: the higher the level of income, the higher the interest rate.

drawn for a *given* supply of real money balances. If real money balances change—for example, if the Fed alters the money supply—the LM curve shifts.

We can use the theory of liquidity preference to understand how monetary policy shifts the LM curve. Suppose that the Fed decreases the money supply from M_1 to M_2 , which causes the supply of real money balances to fall from M_1/P to M_2/P . Figure 10-13 shows what happens. Holding constant the amount of income and thus the demand curve for real money balances, we see that a reduction in the supply of real money balances raises the interest rate that equilibrates the money market. Hence, a decrease in the money supply shifts the LM curve upward.

In summary, the LM curve shows the combinations of the interest rate and the level of income that are consistent with equilibrium in the market for real money balances. The LM curve is drawn for a given supply of real money balances. Decreases in the supply of real money balances shift the LM curve upward. Increases in the supply of real money balances shift the LM curve downward.


A Quantity-Equation Interpretation of the LM Curve

When we first discussed aggregate demand and the short-run determination of income in Chapter 9, we derived the aggregate demand curve from the quantity theory of money. We described the money market with the quantity equation,

$$MV = PY,$$

and assumed that velocity V is constant. This assumption implies that, for any given price level P , the supply of money M by itself determines the level of

figure 10-13



A Reduction in the Money Supply Shifts the LM Curve Upward Panel (a) shows that for any given level of income $\bar{Y}$, a reduction in the money supply raises the interest rate that equilibrates the money market. Therefore, the LM curve in panel (b) shifts upward.

income Y . Because the level of income does not depend on the interest rate, the quantity theory is equivalent to a vertical LM curve.

We can derive the more realistic upward-sloping LM curve from the quantity equation by relaxing the assumption that velocity is constant. The assumption of constant velocity is based on the assumption that the demand for real money balances depends only on the level of income. Yet, as we have noted in our discussion of the liquidity-preference model, the demand for real money balances also depends on the interest rate: a higher interest rate raises the cost of holding money and reduces money demand. When people respond to a higher interest rate by holding less money, each dollar they do hold must be used more often to support a given volume of transactions—that is, the velocity of money must increase. We can write this as

$$MV(r) = PY.$$

The velocity function $V(r)$ indicates that velocity is positively related to the interest rate.

This form of the quantity equation yields an LM curve that slopes upward. Because an increase in the interest rate raises the velocity of money, it raises the level of income for any given money supply and price level. The LM curve expresses this positive relationship between the interest rate and income.

This equation also shows why changes in the money supply shift the LM curve. For any given interest rate and price level, the money supply and the level of income must move together. Thus, increases in the money supply shift the LM curve to the right, and decreases in the money supply shift the LM curve to the left.

Keep in mind that the quantity equation is merely another way to express the theory behind the LM curve. This quantity-theory interpretation of the LM curve is substantively the same as that provided by the theory of liquidity preference. In both cases, the LM curve represents a positive relationship between income and the interest rate that arises from the money market.

Finally, remember that the LM curve by itself does not determine either income Y or the interest rate r that will prevail in the economy. Like the IS curve, the LM curve is only a relationship between these two endogenous variables. The IS and LM curves together determine the economy's equilibrium.


10-3 Conclusion: The Short-Run Equilibrium

We now have all the pieces of the IS – LM model. The two equations of this model are

$$\begin{aligned} Y &= C(Y - T) + I(r) + G && IS, \\ M/P &= L(r, Y) && LM. \end{aligned}$$

The model takes fiscal policy, G and T , monetary policy M , and the price level P as exogenous. Given these exogenous variables, the IS curve provides the combinations of r and Y that satisfy the equation representing the goods market, and the LM curve provides the combinations of r and Y that satisfy the equation representing the money market. These two curves are shown together in Figure 10-14.

figure 10-14




Equilibrium in the IS – LM Model

The intersection of the IS and LM curves represents simultaneous equilibrium in the market for goods and services and in the market for real money balances for given values of government spending, taxes, the money supply, and the price level.

The equilibrium of the economy is the point at which the *IS* curve and the *LM* curve cross. This point gives the interest rate r and the level of income Y that satisfy conditions for equilibrium in both the goods market and the money market. In other words, at this intersection, actual expenditure equals planned expenditure, and the demand for real money balances equals the supply.

As we conclude this chapter, let's recall that our ultimate goal in developing the *IS–LM* model is to analyze short-run fluctuations in economic activity. Figure 10-15 illustrates how the different pieces of our theory fit together. In this chapter we developed the Keynesian cross and the theory of liquidity preference as building blocks for the *IS–LM* model. As we see more fully in the next chapter, the *IS–LM* model helps explain the position and slope of the aggregate demand curve. The aggregate demand curve, in turn, is a piece of the model of aggregate supply and aggregate demand, which economists use to explain the short-run effects of policy changes and other events on national income.

figure 10-15



The Theory of Short-Run Fluctuations This schematic diagram shows how the different pieces of the theory of short-run fluctuations fit together. The Keynesian cross explains the *IS* curve, and the theory of liquidity preference explains the *LM* curve. The *IS* and *LM* curves together yield the *IS–LM* model, which explains the aggregate demand curve. The aggregate demand curve is part of the model of aggregate supply and aggregate demand, which economists use to explain short-run fluctuations in economic activity.

Summary

1. The Keynesian cross is a basic model of income determination. It takes fiscal policy and planned investment as exogenous and then shows that there is one level of national income at which actual expenditure equals planned expenditure. It shows that changes in fiscal policy have a multiplied impact on income.

2. Once we allow planned investment to depend on the interest rate, the Keynesian cross yields a relationship between the interest rate and national income. A higher interest rate lowers planned investment, and this in turn lowers national income. The downward-sloping *IS* curve summarizes this negative relationship between the interest rate and income.
3. The theory of liquidity preference is a basic model of the determination of the interest rate. It takes the money supply and the price level as exogenous and assumes that the interest rate adjusts to equilibrate the supply and demand for real money balances. The theory implies that increases in the money supply lower the interest rate.
4. Once we allow the demand for real money balances to depend on national income, the theory of liquidity preference yields a relationship between income and the interest rate. A higher level of income raises the demand for real money balances, and this in turn raises the interest rate. The upward-sloping *LM* curve summarizes this positive relationship between income and the interest rate.
5. The *IS–LM* model combines the elements of the Keynesian cross and the elements of the theory of liquidity preference. The *IS* curve shows the points that satisfy equilibrium in the goods market, and the *LM* curve shows the points that satisfy equilibrium in the money market. The intersection of the *IS* and *LM* curves shows the interest rate and income that satisfy equilibrium in both markets.

KEY CONCEPTS

<i>IS–LM</i> model	Keynesian cross	Tax multiplier
<i>IS</i> curve	Government-purchases multiplier	Theory of liquidity preference
<i>LM</i> curve		

QUESTIONS FOR REVIEW

1. Use the Keynesian cross to explain why fiscal policy has a multiplied effect on national income.
2. Use the theory of liquidity preference to explain why an increase in the money supply lowers the interest rate. What does this explanation assume about the price level?
3. Why does the *IS* curve slope downward?
4. Why does the *LM* curve slope upward?

PROBLEMS AND APPLICATIONS

1. Use the Keynesian cross to predict the impact of
 - a. An increase in government purchases.
 - b. An increase in taxes.
 - c. An equal increase in government purchases and taxes.
2. In the Keynesian cross, assume that the consumption function is given by

$$C = 200 + 0.75(Y - T).$$
 Planned investment is 100; government purchases and taxes are both 100.

- a. Graph planned expenditure as a function of income.
 - b. What is the equilibrium level of income?
 - c. If government purchases increase to 125, what is the new equilibrium income?
 - d. What level of government purchases is needed to achieve an income of 1,600?
3. Although our development of the Keynesian cross in this chapter assumes that taxes are a fixed amount, in many countries (including the United States) taxes depend on income. Let's represent the tax system by writing tax revenue as

$$T = \bar{T} + tY,$$

where $\bar{T}$ and t are parameters of the tax code. The parameter t is the marginal tax rate: if income rises by \$1, taxes rise by $t \times \$1$.

- a. How does this tax system change the way consumption responds to changes in GDP?
 - b. In the Keynesian cross, how does this tax system alter the government-purchases multiplier?
 - c. In the $IS-LM$ model, how does this tax system alter the slope of the IS curve?
4. Consider the impact of an increase in thriftiness in the Keynesian cross. Suppose the consumption

function is

$$C = \bar{C} + c(Y - T),$$

where $\bar{C}$ is a parameter called *autonomous consumption* and c is the marginal propensity to consume.

- a. What happens to equilibrium income when the society becomes more thrifty, as represented by a decline in $\bar{C}$?
 - b. What happens to equilibrium saving?
 - c. Why do you suppose this result is called the *paradox of thrift*?
 - d. Does this paradox arise in the classical model of Chapter 3? Why or why not?
5. Suppose that the money demand function is

$$(M/P)^d = 1,000 - 100r,$$

where r is the interest rate in percent. The money supply M is 1,000 and the price level P is 2.

- a. Graph the supply and demand for real money balances.
- b. What is the equilibrium interest rate?
- c. Assume that the price level is fixed. What happens to the equilibrium interest rate if the supply of money is raised from 1,000 to 1,200?
- d. If the Fed wishes to raise the interest rate to 7 percent, what money supply should it set?

Aggregate Demand II

Science is a parasite: the greater the patient population the better the advance in physiology and pathology; and out of pathology arises therapy. The year 1932 was the trough of the great depression, and from its rotten soil was belatedly begot a new subject that today we call macroeconomics.

— Paul Samuelson

In Chapter 10 we assembled the pieces of the $IS-LM$ model. We saw that the IS curve represents the equilibrium in the market for goods and services, that the LM curve represents the equilibrium in the market for real money balances, and that the IS and LM curves together determine the interest rate and national income in the short run when the price level is fixed. Now we turn our attention to applying the $IS-LM$ model to analyze three issues.

First, we examine the potential causes of fluctuations in national income. We use the $IS-LM$ model to see how changes in the exogenous variables (government purchases, taxes, and the money supply) influence the endogenous variables (the interest rate and national income). We also examine how various shocks to the goods markets (the IS curve) and the money market (the LM curve) affect the interest rate and national income in the short run.

Second, we discuss how the $IS-LM$ model fits into the model of aggregate supply and aggregate demand we introduced in Chapter 9. In particular, we examine how the $IS-LM$ model provides a theory of the slope and position of the aggregate demand curve. Here we relax the assumption that the price level is fixed, and we show that the $IS-LM$ model implies a negative relationship between the price level and national income. The model can also tell us what events shift the aggregate demand curve and in what direction.

Third, we examine the Great Depression of the 1930s. As this chapter's opening quotation indicates, this episode gave birth to short-run macroeconomic theory, for it led Keynes and his many followers to think that aggregate demand was the key to understanding fluctuations in national income. With the benefit of hindsight, we can use the $IS-LM$ model to discuss the various explanations of this traumatic economic downturn.

11-1 Explaining Fluctuations With the *IS-LM* Model

The intersection of the *IS* curve and the *LM* curve determines the level of national income. When one of these curves shifts, the short-run equilibrium of the economy changes, and national income fluctuates. In this section we examine how changes in policy and shocks to the economy can cause these curves to shift.


How Fiscal Policy Shifts the *IS* Curve and Changes the Short-Run Equilibrium

We begin by examining how changes in fiscal policy (government purchases and taxes) alter the economy's short-run equilibrium. Recall that changes in fiscal policy influence planned expenditure and thereby shift the *IS* curve. The *IS-LM* model shows how these shifts in the *IS* curve affect income and the interest rate.

Changes in Government Purchases Consider an increase in government purchases of ΔG . The government-purchases multiplier in the Keynesian cross tells us that, at any given interest rate, this change in fiscal policy raises the level of income by $\Delta G/(1 - MPC)$. Therefore, as Figure 11-1 shows, the *IS* curve shifts to the right by this amount. The equilibrium of the economy moves from point A to point B. The increase in government purchases raises both income and the interest rate.

To understand fully what's happening in Figure 11-1, it helps to keep in mind the building blocks for the *IS-LM* model from the preceding chapter—the Keynesian cross and the theory of liquidity preference. Here is the story. When the government

figure 11-1



An Increase in Government Purchases in the *IS-LM* Model

An increase in government purchases shifts the *IS* curve to the right. The equilibrium moves from point A to point B. Income rises from Y_1 to Y_2 , and the interest rate rises from r_1 to r_2 .


increases its purchases of goods and services, the economy's planned expenditure rises. The increase in planned expenditure stimulates the production of goods and services, which causes total income Y to rise. These effects should be familiar from the Keynesian cross.

Now consider the money market, as described by the theory of liquidity preference. Because the economy's demand for money depends on income, the rise in total income increases the quantity of money demanded at every interest rate. The supply of money has not changed, however, so higher money demand causes the equilibrium interest rate r to rise.

The higher interest rate arising in the money market, in turn, has ramifications back in the goods market. When the interest rate rises, firms cut back on their investment plans. This fall in investment partially offsets the expansionary effect of the increase in government purchases. Thus, the increase in income in response to a fiscal expansion is smaller in the $IS-LM$ model than it is in the Keynesian cross (where investment is assumed to be fixed). You can see this in Figure 11-1. The horizontal shift in the IS curve equals the rise in equilibrium income in the Keynesian cross. This amount is larger than the increase in equilibrium income here in the $IS-LM$ model. The difference is explained by the crowding out of investment caused by a higher interest rate.

Changes in Taxes In the $IS-LM$ model, changes in taxes affect the economy much the same as changes in government purchases do, except that taxes affect expenditure through consumption. Consider, for instance, a decrease in taxes of ΔT . The tax cut encourages consumers to spend more and, therefore, increases planned expenditure. The tax multiplier in the Keynesian cross tells us that, at any given interest rate, this change in policy raises the level of income by $\Delta T \times MPC / (1 - MPC)$. Therefore, as Figure 11-2 illustrates, the IS curve shifts to the

figure 11-2



A Decrease in Taxes in the $IS-LM$ Model A decrease in taxes shifts the IS curve to the right. The equilibrium moves from point A to point B. Income rises from Y_1 to Y_2 , and the interest rate rises from r_1 to r_2 .

right by this amount. The equilibrium of the economy moves from point A to point B. The tax cut raises both income and the interest rate. Once again, because the higher interest rate depresses investment, the increase in income is smaller in the *IS–LM* model than it is in the Keynesian cross.


How Monetary Policy Shifts the *LM* Curve and Changes the Short-Run Equilibrium

We now examine the effects of monetary policy. Recall that a change in the money supply alters the interest rate that equilibrates the money market for any given level of income and, thereby, shifts the *LM* curve. The *IS–LM* model shows how a shift in the *LM* curve affects income and the interest rate.

Consider an increase in the money supply. An increase in M leads to an increase in real money balances M/P , because the price level P is fixed in the short run. The theory of liquidity preference shows that for any given level of income, an increase in real money balances leads to a lower interest rate. Therefore, the *LM* curve shifts downward, as in Figure 11-3. The equilibrium moves from point A to point B. The increase in the money supply lowers the interest rate and raises the level of income.

Once again, to tell the story that explains the economy's adjustment from point A to point B, we rely on the building blocks of the *IS–LM* model—the Keynesian cross and the theory of liquidity preference. This time, we begin with the money market, where the monetary-policy action occurs. When the Federal Reserve increases the supply of money, people have more money than they want to hold at the prevailing interest rate. As a result, they start depositing this extra money in banks or use it to buy bonds. The interest rate r then falls until people are willing to hold all the extra money that the Fed has created; this brings the money market to a new equilibrium. The lower interest rate, in turn,

figure 11-3



An Increase in the Money Supply in the *IS–LM* Model

An increase in the money supply shifts the *LM* curve downward. The equilibrium moves from point A to point B. Income rises from Y_1 to Y_2 , and the interest rate falls from r_1 to r_2 .

has ramifications for the goods market. A lower interest rate stimulates planned investment, which increases planned expenditure, production, and income Y .

Thus, the $IS-LM$ model shows that monetary policy influences income by changing the interest rate. This conclusion sheds light on our analysis of monetary policy in Chapter 9. In that chapter we showed that in the short run, when prices are sticky, an expansion in the money supply raises income. But we did not discuss *how* a monetary expansion induces greater spending on goods and services—a process called the **monetary transmission mechanism**. The $IS-LM$ model shows that an increase in the money supply lowers the interest rate, which stimulates investment and thereby expands the demand for goods and services.

The Interaction Between Monetary and Fiscal Policy

When analyzing any change in monetary or fiscal policy, it is important to keep in mind that the policymakers who control these policy tools are aware of what the other policymakers are doing. A change in one policy, therefore, may influence the other, and this interdependence may alter the impact of a policy change.

For example, suppose Congress were to raise taxes. What effect should this policy have on the economy? According to the $IS-LM$ model, the answer depends on how the Fed responds to the tax increase.

Figure 11-4 shows three of the many possible outcomes. In panel (a), the Fed holds the money supply constant. The tax increase shifts the IS curve to the left. Income falls (because higher taxes reduce consumer spending), and the interest rate falls (because lower income reduces the demand for money). The fall in income indicates that the tax hike causes a recession.


In panel (b), the Fed wants to hold the interest rate constant. In this case, when the tax increase shifts the IS curve to the left, the Fed must decrease the money supply to keep the interest rate at its original level. This fall in the money supply shifts the LM curve upward. The interest rate does not fall, but income falls by a larger amount than if the Fed had held the money supply constant. Whereas in panel (a) the lower interest rate stimulated investment and partially offset the contractionary effect of the tax hike, in panel (b) the Fed deepens the recession by keeping the interest rate high.

In panel (c), the Fed wants to prevent the tax increase from lowering income. It must, therefore, raise the money supply and shift the LM curve downward enough to offset the shift in the IS curve. In this case, the tax increase does not cause a recession, but it does cause a large fall in the interest rate. Although the level of income is not changed, the combination of a tax increase and a monetary expansion does change the allocation of the economy's resources. The higher taxes depress consumption, while the lower interest rate stimulates investment. Income is not affected because these two effects exactly balance.


From this example we can see that the impact of a change in fiscal policy depends on the policy the Fed pursues—that is, on whether it holds the money supply, the interest rate, or the level of income constant. More generally, whenever analyzing a change in one policy, we must make an assumption about its effect on the other policy. The most appropriate assumption depends on the case at hand and the many political considerations that lie behind economic policymaking.

figure 11-4


(a) Fed Holds Money Supply Constant



(b) Fed Holds Interest Rate Constant



(c) Fed Holds Income Constant



The Response of the Economy to a Tax Increase

How the economy responds to a tax increase depends on how the monetary authority responds. In panel (a) the Fed holds the money supply constant. In panel (b) the Fed holds the interest rate constant by reducing the money supply. In panel (c) the Fed holds the level of income constant by raising the money supply.

CASE STUDY

Policy Analysis With Macroeconometric Models

The *IS–LM* model shows how monetary and fiscal policy influence the equilibrium level of income. The predictions of the model, however, are qualitative, not quantitative. The *IS–LM* model shows that increases in government purchases raise GDP and that increases in taxes lower GDP. But when economists analyze specific policy proposals, they need to know not only the direction of the effect but also the size. For example, if Congress increases taxes by \$100 billion and if monetary policy is not altered, how much will GDP fall? To answer this question, economists need to go beyond the graphical representation of the *IS–LM* model.

Macroeconometric models of the economy provide one way to evaluate policy proposals. A *macroeconometric model* is a model that describes the economy quantitatively, rather than only qualitatively. Many of these models are essentially more complicated and more realistic versions of our *IS–LM* model. The economists who build macroeconometric models use historical data to estimate parameters such as the marginal propensity to consume, the sensitivity of investment to the interest rate, and the sensitivity of money demand to the interest rate. Once a model is built, economists can simulate the effects of alternative policies with the help of a computer.

Table 11-1 shows the fiscal-policy multipliers implied by one widely used macroeconometric model, the Data Resources Incorporated (DRI) model, named for the economic forecasting firm that developed it. The multipliers are given for two assumptions about how the Fed might respond to changes in fiscal policy.

One assumption about monetary policy is that the Fed keeps the nominal interest rate constant. That is, when fiscal policy shifts the *IS* curve to the right or to the left, the Fed adjusts the money supply to shift the *LM* curve in the same direction. Because there is no crowding out of investment due to a changing interest rate, the fiscal-policy multipliers are similar to those from the Keynesian cross. The DRI model indicates that, in this case, the government-purchases multiplier is 1.93, and the tax multiplier is -1.19 . That is, a \$100 billion increase in government purchases raises GDP by \$193 billion, and a \$100 billion increase in taxes lowers GDP by \$119 billion.

table 11-1

The Fiscal-Policy Multipliers in the DRI Model

Assumption About Monetary Policy	VALUE OF MULTIPLIERS	
	$\Delta Y / \Delta G$	$\Delta Y / \Delta T$
Nominal interest rate held constant	1.93	-1.19
Money supply held constant	0.60	-0.26

Note: This table gives the fiscal-policy multipliers for a sustained change in government purchases or in personal income taxes. These multipliers are for the fourth quarter after the policy change is made.

Source: Otto Eckstein, *The DRI Model of the U.S. Economy* (New York: McGraw-Hill, 1983), 169.

The second assumption about monetary policy is that the Fed keeps the money supply constant so that the LM curve does not shift. In this case, the interest rate rises, and investment is crowded out, so the multipliers are much smaller. The government-purchases multiplier is only 0.60, and the tax multiplier is only -0.26 . That is, a \$100 billion increase in government purchases raises GDP by \$60 billion, and a \$100 billion increase in taxes lowers GDP by \$26 billion.

Table 11-1 shows that the fiscal-policy multipliers are very different under the two assumptions about monetary policy. The impact of any change in fiscal policy depends crucially on how the Fed responds to that change.

Shocks in the $IS-LM$ Model

Because the $IS-LM$ model shows how national income is determined in the short run, we can use the model to examine how various economic disturbances affect income. So far we have seen how changes in fiscal policy shift the IS curve and how changes in monetary policy shift the LM curve. Similarly, we can group other disturbances into two categories: shocks to the IS curve and shocks to the LM curve.

Shocks to the IS curve are exogenous changes in the demand for goods and services. Some economists, including Keynes, have emphasized that such changes in demand can arise from investors' *animal spirits*—exogenous and perhaps self-fulfilling waves of optimism and pessimism. For example, suppose that firms become pessimistic about the future of the economy and that this pessimism causes them to build fewer new factories. This reduction in the demand for investment goods causes a contractionary shift in the investment function: at every interest rate, firms want to invest less. The fall in investment reduces planned expenditure and shifts the IS curve to the left, reducing income and employment. This fall in equilibrium income in part validates the firms' initial pessimism.

Shocks to the IS curve may also arise from changes in the demand for consumer goods. Suppose, for instance, that the election of a popular president increases consumer confidence in the economy. This induces consumers to save less for the future and consume more today. We can interpret this change as an upward shift in the consumption function. This shift in the consumption function increases planned expenditure and shifts the IS curve to the right, and this raises income.



Calvin and Hobbes © 1992 Watterson.
Dist. by Universal Press Syndicate.

Shocks to the LM curve arise from exogenous changes in the demand for money. For example, suppose that new restrictions on credit-card availability increase the amount of money people choose to hold. According to the theory of liquidity preference, when money demand rises, the interest rate necessary to equilibrate the money market is higher (for any given level of income and money supply). Hence, an increase in money demand shifts the LM curve upward, which tends to raise the interest rate and depress income.

In summary, several kinds of events can cause economic fluctuations by shifting the IS curve or the LM curve. Remember, however, that such fluctuations are not inevitable. Policymakers can try to use the tools of monetary and fiscal policy to offset exogenous shocks. If policymakers are sufficiently quick and skillful (admittedly, a big if), shocks to the IS or LM curves need not lead to fluctuations in income or employment.

CASE STUDY

The U.S. Slowdown of 2001

In 2001, the U.S. economy experienced a pronounced slowdown in economic activity. The unemployment rate rose from 3.9 percent in October 2000 to 4.9 percent in August 2001, and then to 5.8 percent in December 2001. In many ways, the slowdown looked like a typical recession driven by a fall in aggregate demand.

Two notable shocks can help explain this event. The first was a decline in the stock market. During the 1990s, the stock market experienced a boom of historic proportions, as investors became optimistic about the prospects of the new information technology. Some economists viewed the optimism as excessive at the time, and in hindsight this proved to be the case. When the optimism faded, average stock prices fell by about 25 percent from August 2000 to August 2001. The fall in the market reduced household wealth and thus consumer spending. In addition, the declining perceptions of the profitability of the new technologies led to a fall in investment spending. In the language of the $IS-LM$ model, the IS curve shifted to the left.

The second shock was the terrorist attacks on New York and Washington on September 11, 2001. In the week after the attacks, the stock market fell another 12 percent, its biggest weekly loss since the Great Depression of the 1930s. Moreover, the attacks increased uncertainty about what the future would hold. Uncertainty can reduce spending because households and firms postpone some of their plans until the uncertainty is resolved. Thus, the terrorist attacks shifted the IS curve further to the left.

Fiscal and monetary policymakers were quick to respond to these events. Congress passed a tax cut in 2001, including an immediate tax rebate. One goal of the tax cut was to stimulate consumer spending. After the terrorist attacks, Congress increased government spending by appropriating funds to rebuild New York and to bail out the ailing airline industry. Both of these fiscal measures shifted the IS curve to the right.

At the same time, the Fed pursued expansionary monetary policy, shifting the LM curve to the right. Money growth accelerated, and interest rates fell. The interest rate on three-month Treasury bills fell from 6.4 percent in November of 2000 to 3.3 percent in August 2001, and then to 2.1 percent in September 2001 in the immediate aftermath of the terrorist attacks.

The magnitude of the slowdown of 2001 was not yet determined as this book was going to press. The big question was whether the policy measures undertaken were sufficient to offset the shocks that the economy had suffered. By the time you are reading this, you may know the answer.

What Is the Fed's Policy Instrument—The Money Supply or the Interest Rate?

Our analysis of monetary policy has been based on the assumption that the Fed influences the economy by controlling the money supply. By contrast, when the media report on changes in Fed policy, they often simply say that the Fed has raised or lowered interest rates. Which is right? Even though these two views may seem different, both are correct, and it is important to understand why.

In recent years, the Fed has used the *federal funds rate*—the interest rate that banks charge one another for overnight loans—as its short-term policy instrument. When the Federal Open Market Committee meets every six weeks to set monetary policy, it votes on a target for this interest rate that will apply until the next meeting. After the meeting is over, the Fed's bond traders in New York are told to conduct the open-market operations necessary to hit that target. These open-market operations change the money supply and shift the LM curve so that the equilibrium interest rate (determined by the intersection of the IS and LM curves) equals the target interest rate that the Federal Open Market Committee has chosen.

As a result of this operating procedure, Fed policy is often discussed in terms of changing interest rates. Keep in mind, however, that behind these changes in interest rates are the necessary changes in the money supply. A newspaper might report, for instance, that “the Fed has lowered interest rates.” To be more precise, we can translate this statement as meaning “the Federal Open Market Committee has instructed the Fed bond traders to buy bonds in open-market operations so as to increase the money supply, shift the LM curve, and reduce the equilibrium interest rate to hit a new lower target.”

Why has the Fed chosen to use an interest rate, rather than the money supply, as its short-term policy instrument? One possible answer is that shocks to the LM curve are more prevalent than shocks to the IS curve. When the Fed targets interest rates, it automatically offsets LM shocks by altering the money supply, but the policy exacerbates IS shocks. If LM shocks are the more prevalent type, then a policy of targeting the interest rate leads to greater economic stability than a policy of targeting the money supply. (Problem 7 at the end of this chapter asks you to analyze this issue more fully.)

Another possible reason for using the interest rate as the short-term policy instrument is that interest rates are easier to measure than the money supply. As we

saw in Chapter 4, the Fed has several different measures of money— $M1$, $M2$, and so on—which sometimes move in different directions. Rather than deciding which measure is best, the Fed avoids the question by using the federal funds rate as its policy instrument.

11-2 *IS-LM* as a Theory of Aggregate Demand

We have been using the *IS-LM* model to explain national income in the short run when the price level is fixed. To see how the *IS-LM* model fits into the model of aggregate supply and aggregate demand introduced in Chapter 9, we now examine what happens in the *IS-LM* model if the price level is allowed to change. As was promised when we began our study of this model, the *IS-LM* model provides a theory to explain the position and slope of the aggregate demand curve.

From the *IS-LM* Model to the Aggregate Demand Curve


Recall from Chapter 9 that the aggregate demand curve describes a relationship between the price level and the level of national income. In Chapter 9 this relationship was derived from the quantity theory of money. The analysis showed that for a given money supply, a higher price level implies a lower level of income. Increases in the money supply shift the aggregate demand curve to the right, and decreases in the money supply shift the aggregate demand curve to the left.

To understand the determinants of aggregate demand more fully, we now use the *IS-LM* model, rather than the quantity theory, to derive the aggregate demand curve. First, we use the *IS-LM* model to show why national income falls as the price level rises—that is, why the aggregate demand curve is downward sloping. Second, we examine what causes the aggregate demand curve to shift.

To explain why the aggregate demand curve slopes downward, we examine what happens in the *IS-LM* model when the price level changes. This is done in Figure 11-5. For any given money supply M , a higher price level P reduces the supply of real money balances M/P . A lower supply of real money balances shifts the *LM* curve upward, which raises the equilibrium interest rate and lowers the equilibrium level of income, as shown in panel (a). Here the price level rises from P_1 to P_2 , and income falls from Y_1 to Y_2 . The aggregate demand curve in panel (b) plots this negative relationship between national income and the price level. In other words, the aggregate demand curve shows the set of equilibrium points that arise in the *IS-LM* model as we vary the price level and see what happens to income.

What causes the aggregate demand curve to shift? Because the aggregate demand curve is merely a summary of results from the *IS-LM* model, events that shift the *IS* curve or the *LM* curve (for a given price level) cause the aggregate demand curve to shift. For instance, an increase in the money supply raises income in the *IS-LM* model for any given price level; it thus shifts the

figure 11-5



Deriving the Aggregate Demand Curve With the IS-LM Model Panel (a) shows the IS-LM model: an increase in the price level from P_1 to P_2 lowers real money balances and thus shifts the LM curve upward. The shift in the LM curve lowers income from Y_1 to Y_2 . Panel (b) shows the aggregate demand curve summarizing this relationship between the price level and income: the higher the price level, the lower the level of income.


aggregate demand curve to the right, as shown in panel (a) of Figure 11-6. Similarly, an increase in government purchases or a decrease in taxes raises income in the IS-LM model for a given price level; it also shifts the aggregate demand curve to the right, as shown in panel (b) of Figure 11-6. Conversely, a decrease in the money supply, a decrease in government purchases, or an increase in taxes lowers income in the IS-LM model and shifts the aggregate demand curve to the left.

We can summarize these results as follows: *A change in income in the IS-LM model resulting from a change in the price level represents a movement along the aggregate demand curve. A change in income in the IS-LM model for a fixed price level represents a shift in the aggregate demand curve.*

The IS-LM Model in the Short Run and Long Run

The IS-LM model is designed to explain the economy in the short run when the price level is fixed. Yet, now that we have seen how a change in the price level influences the equilibrium in the IS-LM model, we can also use the model to describe the economy in the long run when the price level adjusts to ensure that the economy produces at its natural rate. By using the IS-LM model to describe the long run, we can show clearly how the Keynesian model of income determination differs from the classical model of Chapter 3.


figure 11-6



How Monetary and Fiscal Policies Shift the Aggregate Demand Curve Panel (a) shows a monetary expansion. For any given price level, an increase in the money supply raises real money balances, shifts the LM curve downward, and raises income. Hence, an increase in the money supply shifts the aggregate demand curve to the right. Panel (b) shows a fiscal expansion, such as an increase in government purchases or a decrease in taxes. The fiscal expansion shifts the IS curve to the right and, for any given price level, raises income. Hence, a fiscal expansion shifts the aggregate demand curve to the right.

Panel (a) of Figure 11-7 shows the three curves that are necessary for understanding the short-run and long-run equilibria: the IS curve, the LM curve, and the vertical line representing the natural rate of output $\bar{Y}$. The LM curve is, as always, drawn for a fixed price level, P_1 . The short-run equilibrium of the economy is point K, where the IS curve crosses the LM curve. Notice that in this short-run equilibrium, the economy's income is less than its natural rate.

figure 11-7



The Short-Run and Long-Run Equilibria We can compare the short-run and long-run equilibria using either the *IS*–*LM* diagram in panel (a) or the aggregate supply–aggregate demand diagram in panel (b). In the short run, the price level is stuck at P_1 . The short-run equilibrium of the economy is therefore point K. In the long run, the price level adjusts so that the economy is at the natural rate. The long-run equilibrium is therefore point C.

Panel (b) of Figure 11-7 shows the same situation in the diagram of aggregate supply and aggregate demand. At the price level P_1 , the quantity of output demanded is below the natural rate. In other words, at the existing price level, there is insufficient demand for goods and services to keep the economy producing at its potential.

In these two diagrams we can examine the short-run equilibrium at which the economy finds itself and the long-run equilibrium toward which the economy gravitates. Point K describes the short-run equilibrium, because it assumes that the price level is stuck at P_1 . Eventually, the low demand for goods and services causes prices to fall, and the economy moves back toward its natural rate. When the price level reaches P_2 , the economy is at point C, the long-run equilibrium. The diagram of aggregate supply and aggregate demand shows that at point C, the quantity of goods and services demanded equals the natural rate of output. This long-run equilibrium is achieved in the *IS*–*LM* diagram by a shift in the *LM* curve: the fall in the price level raises real money balances and therefore shifts the *LM* curve to the right.

We can now see the key difference between Keynesian and classical approaches to the determination of national income. The Keynesian assumption (represented by point K) is that the price level is stuck. Depending on monetary policy, fiscal policy, and the other determinants of aggregate demand, output may deviate from the natural rate. The classical assumption (represented by point C) is that the price level is fully flexible. The price level adjusts to ensure that national income is always at the natural rate.

To make the same point somewhat differently, we can think of the economy as being described by three equations. The first two are the *IS* and *LM* equations:

$$\begin{aligned} Y &= C(Y - T) + I(r) + G && IS, \\ M/P &= L(r, Y) && LM. \end{aligned}$$

The *IS* equation describes the goods market, and the *LM* equation describes the money market. These two equations contain three endogenous variables: Y , P , and r . The Keynesian approach is to complete the model with the assumption of fixed prices, so the Keynesian third equation is

$$P = P_1.$$

This assumption implies that r and Y must adjust to satisfy the *IS* and *LM* equations. The classical approach is to complete the model with the assumption that output reaches the natural rate, so the classical third equation is

$$Y = \bar{Y}.$$

This assumption implies that r and P must adjust to satisfy the *IS* and *LM* equations.

Which assumption is most appropriate? The answer depends on the time horizon. The classical assumption best describes the long run. Hence, our long-run analysis of national income in Chapter 3 and prices in Chapter 4 assumes that output equals the natural rate. The Keynesian assumption best describes the short run. Therefore, our analysis of economic fluctuations relies on the assumption of a fixed price level.

11-3 The Great Depression

Now that we have developed the model of aggregate demand, let's use it to address the question that originally motivated Keynes: What caused the Great Depression? Even today, more than half a century after the event, economists continue to debate the cause of this major economic downturn. The Great Depression provides an extended case study to show how economists use the *IS-LM* model to analyze economic fluctuations.¹

Before turning to the explanations economists have proposed, look at Table 11-2, which presents some statistics regarding the Depression. These statistics are the battlefield on which debate about the Depression takes place. What do you think happened? An *IS* shift? An *LM* shift? Or something else?

¹ For a flavor of the debate, see Milton Friedman and Anna J. Schwartz, *A Monetary History of the United States, 1867–1960* (Princeton, NJ: Princeton University Press, 1963); Peter Temin, *Did Monetary Forces Cause the Great Depression?* (New York: W.W. Norton, 1976); the essays in Karl Brunner, ed., *The Great Depression Revisited* (Boston: Martinus Nijhoff, 1981); and the symposium on the Great Depression in the Spring 1993 issue of the *Journal of Economic Perspectives*.

table 11-2

What Happened During the Great Depression?

Year	Unemployment Rate (1)	Real GNP (2)	Consumption (2)	Investment (2)	Government Purchases (2)
1929	3.2	203.6	139.6	40.4	22.0
1930	8.9	183.5	130.4	27.4	24.3
1931	16.3	169.5	126.1	16.8	25.4
1932	24.1	144.2	114.8	4.7	24.2
1933	25.2	141.5	112.8	5.3	23.3
1934	22.0	154.3	118.1	9.4	26.6
1935	20.3	169.5	125.5	18.0	27.0
1936	17.0	193.2	138.4	24.0	31.8
1937	14.3	203.2	143.1	29.9	30.8
1938	19.1	192.9	140.2	17.0	33.9
1939	17.2	209.4	148.2	24.7	35.2
1940	14.6	227.2	155.7	33.0	36.4

Source: *Historical Statistics of the United States, Colonial Times to 1970, Parts I and II* (Washington, DC: U.S. Department of Commerce, Bureau of Census, 1975).

Note: (1) The unemployment rate is series D9. (2) Real GNP, consumption, investment, and government purchases are series F3, F48, F52, and F66, and are measured in billions of 1958 dollars. (3) The interest rate is the prime Commercial

The Spending Hypothesis: Shocks to the *IS* Curve

Table 11-2 shows that the decline in income in the early 1930s coincided with falling interest rates. This fact has led some economists to suggest that the cause of the decline may have been a contractionary shift in the *IS* curve. This view is sometimes called the *spending hypothesis*, because it places primary blame for the Depression on an exogenous fall in spending on goods and services.

Economists have attempted to explain this decline in spending in several ways. Some argue that a downward shift in the consumption function caused the contractionary shift in the *IS* curve. The stock market crash of 1929 may have been partly responsible for this shift: by reducing wealth and increasing uncertainty about the future prospects of the U.S. economy, the crash may have induced consumers to save more of their income rather than spending it.

Others explain the decline in spending by pointing to the large drop in investment in housing. Some economists believe that the residential investment boom of the 1920s was excessive and that once this “overbuilding” became apparent, the demand for residential investment declined drastically. Another possible explanation for the fall in residential investment is the reduction in immigration in the 1930s: a more slowly growing population demands less new housing.

Once the Depression began, several events occurred that could have reduced spending further. First, many banks failed in the early 1930s, in part because of inadequate bank regulation, and these bank failures may have exacerbated the fall in investment spending. Banks play the crucial role of getting the funds available

Year	Nominal Interest Rate (3)	Money Supply (4)	Price Level (5)	Inflation (6)	Real Money Balances (7)
1929	5.9	26.6	50.6	—	52.6
1930	3.6	25.8	49.3	−2.6	52.3
1931	2.6	24.1	44.8	−10.1	54.5
1932	2.7	21.1	40.2	−9.3	52.5
1933	1.7	19.9	39.3	−2.2	50.7
1934	1.0	21.9	42.2	7.4	51.8
1935	0.8	25.9	42.6	0.9	60.8
1936	0.8	29.6	42.7	0.2	62.9
1937	0.9	30.9	44.5	4.2	69.5
1938	0.8	30.5	43.9	−1.3	69.5
1939	0.6	34.2	43.2	−1.6	79.1
1940	0.6	39.7	43.9	1.6	90.3

Paper rate, 4–6 months, series X445. (4) The money supply is series X414, currency plus demand deposits, measured in billions of dollars. (5) The price level is the GNP deflator (1958 = 100), series E1. (6) The inflation rate is the percentage change in the price level series. (7) Real money balances, calculated by dividing the money supply by the price level and multiplying by 100, are in billions of 1958 dollars.

for investment to those households and firms that can best use them. The closing of many banks in the early 1930s may have prevented some businesses from getting the funds they needed for capital investment and, therefore, may have led to a further contractionary shift in the investment function.²

In addition, the fiscal policy of the 1930s caused a contractionary shift in the *IS* curve. Politicians at that time were more concerned with balancing the budget than with using fiscal policy to keep production and employment at their natural rates. The Revenue Act of 1932 increased various taxes, especially those falling on lower- and middle-income consumers.³ The Democratic platform of that year expressed concern about the budget deficit and advocated an “immediate and drastic reduction of governmental expenditures.” In the midst of historically high unemployment, policymakers searched for ways to raise taxes and reduce government spending.

There are, therefore, several ways to explain a contractionary shift in the *IS* curve. Keep in mind that these different views may all be true. There may be no single explanation for the decline in spending. It is possible that all of these changes coincided and that together they led to a massive reduction in spending.

² Ben Bernanke, “Non-Monetary Effects of the Financial Crisis in the Propagation of the Great Depression,” *American Economic Review* 73 (June 1983): 257–276.

³ E. Cary Brown, “Fiscal Policy in the ‘Thirties: A Reappraisal,” *American Economic Review* 46 (December 1956): 857–879.

The Money Hypothesis: A Shock to the *LM* Curve

Table 11–2 shows that the money supply fell 25 percent from 1929 to 1933, during which time the unemployment rate rose from 3.2 percent to 25.2 percent. This fact provides the motivation and support for what is called the *money hypothesis*, which places primary blame for the Depression on the Federal Reserve for allowing the money supply to fall by such a large amount.⁴ The best-known advocates of this interpretation are Milton Friedman and Anna Schwartz, who defend it in their treatise on U.S. monetary history. Friedman and Schwartz argue that contractions in the money supply have caused most economic downturns and that the Great Depression is a particularly vivid example.

Using the *IS–LM* model, we might interpret the money hypothesis as explaining the Depression by a contractionary shift in the *LM* curve. Seen in this way, however, the money hypothesis runs into two problems.

The first problem is the behavior of *real* money balances. Monetary policy leads to a contractionary shift in the *LM* curve only if real money balances fall. Yet from 1929 to 1931 real money balances rose slightly, because the fall in the money supply was accompanied by an even greater fall in the price level. Although the monetary contraction may be responsible for the rise in unemployment from 1931 to 1933, when real money balances did fall, it cannot easily explain the initial downturn from 1929 to 1931.

The second problem for the money hypothesis is the behavior of interest rates. If a contractionary shift in the *LM* curve triggered the Depression, we should have observed higher interest rates. Yet nominal interest rates fell continuously from 1929 to 1933.

These two reasons appear sufficient to reject the view that the Depression was instigated by a contractionary shift in the *LM* curve. But was the fall in the money stock irrelevant? Next, we turn to another mechanism through which monetary policy might have been responsible for the severity of the Depression—the deflation of the 1930s.

The Money Hypothesis Again: The Effects of Falling Prices

From 1929 to 1933 the price level fell 25 percent. Many economists blame this deflation for the severity of the Great Depression. They argue that the deflation may have turned what in 1931 was a typical economic downturn into an unprecedented period of high unemployment and depressed income. If it is correct, this argument gives new life to the money hypothesis. Because the falling money supply was, plausibly, responsible for the falling price level, it could have been responsible for the severity of the Depression. To evaluate this argument, we must discuss how changes in the price level affect income in the *IS–LM* model.

⁴ We discuss the reasons for this large decrease in the money supply in Chapter 18, where we examine the money supply process in more detail. In particular, see the case study “Bank Failures and the Money Supply in the 1930s.”

The Stabilizing Effects of Deflation In the *IS–LM* model we have developed so far, falling prices raise income. For any given supply of money M , a lower price level implies higher real money balances M/P . An increase in real money balances causes an expansionary shift in the *LM* curve, which leads to higher income.

Another channel through which falling prices expand income is called the **Pigou effect**. Arthur Pigou, a prominent classical economist in the 1930s, pointed out that real money balances are part of households' wealth. As prices fall and real money balances rise, consumers should feel wealthier and spend more. This increase in consumer spending should cause an expansionary shift in the *IS* curve, also leading to higher income.

These two reasons led some economists in the 1930s to believe that falling prices would help stabilize the economy. That is, they thought that a decline in the price level would automatically push the economy back toward full employment. Yet other economists were less confident in the economy's ability to correct itself. They pointed to other effects of falling prices, to which we now turn.

The Destabilizing Effects of Deflation Economists have proposed two theories to explain how falling prices could depress income rather than raise it. The first, called the **debt-deflation theory**, describes the effects of unexpected falls in the price level. The second explains the effects of expected deflation.

The debt-deflation theory begins with an observation from Chapter 4: unanticipated changes in the price level redistribute wealth between debtors and creditors. If a debtor owes a creditor \$1,000, then the real amount of this debt is $\$1,000/P$, where P is the price level. A fall in the price level raises the real amount of this debt—the amount of purchasing power the debtor must repay the creditor. Therefore, an unexpected deflation enriches creditors and impoverishes debtors.

The debt-deflation theory then posits that this redistribution of wealth affects spending on goods and services. In response to the redistribution from debtors to creditors, debtors spend less and creditors spend more. If these two groups have equal spending propensities, there is no aggregate impact. But it seems reasonable to assume that debtors have higher propensities to spend than creditors—perhaps that is why the debtors are in debt in the first place. In this case, debtors reduce their spending by more than creditors raise theirs. The net effect is a reduction in spending, a contractionary shift in the *IS* curve, and lower national income.

To understand how *expected* changes in prices can affect income, we need to add a new variable to the *IS–LM* model. Our discussion of the model so far has not distinguished between the nominal and real interest rates. Yet we know from previous chapters that investment depends on the real interest rate and that money demand depends on the nominal interest rate. If i is the nominal interest rate and π^e is expected inflation, then the *ex ante* real interest rate is $i - \pi^e$. We can now write the *IS–LM* model as

$$\begin{aligned} Y &= C(Y - T) + I(i - \pi^e) + G && IS, \\ M/P &= L(i, Y) && LM. \end{aligned}$$

Expected inflation enters as a variable in the *IS* curve. Thus, changes in expected inflation shift the *IS* curve.


Let's use this extended *IS–LM* model to examine how changes in expected inflation influence the level of income. We begin by assuming that everyone expects the price level to remain the same. In this case, there is no expected inflation ($\pi^e = 0$), and these two equations produce the familiar *IS–LM* model. Figure 11-8 depicts this initial situation with the *LM* curve and the *IS* curve labeled IS_1 . The intersection of these two curves determines the nominal and real interest rates, which for now are the same.

Now suppose that everyone suddenly expects that the price level will fall in the future, so that π^e becomes negative. The real interest rate is now higher at any given nominal interest rate. This increase in the real interest rate depresses planned investment spending, shifting the *IS* curve from IS_1 to IS_2 . Thus, an expected deflation leads to a reduction in national income from Y_1 to Y_2 . The nominal interest rate falls from i_1 to i_2 , whereas the real interest rate rises from r_1 to r_2 .

Here is the story behind this figure. When firms come to expect deflation, they become reluctant to borrow to buy investment goods because they believe they will have to repay these loans later in more valuable dollars. The fall in investment depresses planned expenditure, which in turn depresses income. The fall in income reduces the demand for money, and this reduces the nominal interest rate that equilibrates the money market. The nominal interest rate falls by less than the expected deflation, so the real interest rate rises.

Note that there is a common thread in these two stories of destabilizing deflation. In both, falling prices depress national income by causing a contractionary shift in the *IS* curve. Because a deflation of the size observed from 1929 to 1933 is unlikely except in the presence of a major contraction in the money supply, these two explanations give some of the responsibility for the Depression—especially its severity—to the Fed. In other words, if falling prices are destabilizing, then a contraction in the money supply can lead to a fall in income, even without a decrease in real money balances or a rise in nominal interest rates.

figure 11-8



Expected Deflation in the *IS–LM* Model An expected deflation (a negative value of π^e) raises the real interest rate for any given nominal interest rate, and this depresses investment spending. The reduction in investment shifts the *IS* curve downward. The level of income falls from Y_1 to Y_2 . The nominal interest rate falls from i_1 to i_2 , and the real interest rate rises from r_1 to r_2 .

Could the Depression Happen Again?

Economists study the Depression both because of its intrinsic interest as a major economic event and to provide guidance to policymakers so that it will not happen again. To state with confidence whether this event could recur, we would need to know why it happened. Because there is not yet agreement on the causes of the Great Depression, it is impossible to rule out with certainty another depression of this magnitude.

Yet most economists believe that the mistakes that led to the Great Depression are unlikely to be repeated. The Fed seems unlikely to allow the money supply to fall by one-fourth. Many economists believe that the deflation of the early 1930s was responsible for the depth and length of the Depression. And it seems likely that such a prolonged deflation was possible only in the presence of a falling money supply.

The fiscal-policy mistakes of the Depression are also unlikely to be repeated. Fiscal policy in the 1930s not only failed to help but actually further depressed aggregate demand. Few economists today would advocate such a rigid adherence to a balanced budget in the face of massive unemployment.

In addition, there are many institutions today that would help prevent the events of the 1930s from recurring. The system of Federal Deposit Insurance makes widespread bank failures less likely. The income tax causes an automatic reduction in taxes when income falls, which stabilizes the economy. Finally, economists know more today than they did in the 1930s. Our knowledge of how the economy works, limited as it still is, should help policymakers formulate better policies to combat such widespread unemployment.

CASE STUDY

The Japanese Slump

During the 1990s, after many years of rapid growth and enviable prosperity, the Japanese economy experienced a prolonged downturn. Real GDP grew at an average rate of only 1.3 percent over the decade, compared with 4.3 percent over the previous twenty years. The unemployment rate, which had historically been very low in Japan, rose from 2.1 percent in 1990 to 4.7 percent in 1999. In August 2001, unemployment hit 5.0 percent, the highest rate since the government began compiling the statistic in 1953.

Although the Japanese slump of the 1990s is not even close in magnitude to the Great Depression of the 1930s, the episodes are similar in several ways. First, both episodes are traced in part to a large decline in stock prices. In Japan, stock prices at the end of the 1990s were less than half the peak level they had reached about a decade earlier. Like the stock market, Japanese land prices had also skyrocketed in the 1980s before crashing in the 1990s. (At the peak of Japan's land bubble, it was said that the land under the Imperial Palace was worth more than the entire state of California.) When stock and land prices collapsed, Japanese citizens saw their wealth plummet. This decline in wealth, like that during the Great Depression, depressed consumer spending.

Second, during both episodes, banks ran into trouble and exacerbated the slump in economic activity. Japanese banks in the 1980s had made many loans that were backed by stock or land. When the value of this collateral fell, borrowers started defaulting on their loans. These defaults on the old loans reduced the banks' ability to make new loans. The resulting "credit crunch" made it harder for firms to finance investment projects and, thus, depressed investment spending.

Third, both episodes saw a fall in economic activity coincide with very low interest rates. In Japan in the 1990s, as in the United States in the 1930s, short-term nominal interest rates were less than 1 percent. This fact suggests that the cause of the slump was primarily a contractionary shift in the *IS* curve, because such a shift reduces both income and the interest rate. The obvious suspects to explain the *IS* shift are the crashes in stock and land prices and the problems in the banking system.

Finally, the policy debate in Japan mirrored the debate over the Great Depression. Some economists recommended that the Japanese government pass large tax cuts to encourage more consumer spending. Although this advice was followed to some extent, Japanese policymakers were reluctant to enact very large tax cuts because, like the U.S. policymakers in the 1930s, they wanted to avoid budget deficits. In Japan, this reluctance to increase government debt arose in part because the government was facing a large unfunded pension liability and a rapidly aging population.

Other economists recommended that the Bank of Japan expand the money supply more rapidly. Even if nominal interest rates could not go much lower, then perhaps more rapid money growth could raise expected inflation, lower real interest rates, and stimulate investment spending. Thus, although economists differed about whether fiscal or monetary policy was more likely to be effective, there was wide agreement that the solution to Japan's slump, like the solution to the Great Depression, rested in more aggressive expansion of aggregate demand.⁵

11-4 Conclusion

The purpose of this chapter and the previous one has been to deepen our understanding of aggregate demand. We now have the tools to analyze the effects of monetary and fiscal policy in the long run and in the short run. In the long run, prices are flexible, and we use the classical analysis of Parts II and III of this book. In the short run, prices are sticky, and we use the *IS-LM* model to examine how changes in policy influence the economy.

⁵ To learn more about this episode, see Adam S. Posen, *Restoring Japan's Economic Growth* (Washington, DC: Institute for International Economics, 1998).



The Liquidity Trap

In Japan in the 1990s and the United States in the 1930s, interest rates reached very low levels. As Table 11-2 shows, U.S. interest rates were well under 1 percent throughout the second half of the 1930s. The same was true in Japan during the second half of the 1990s. In 1999, Japanese short-term interest rates fell to about one-tenth of 1 percent.

Some economists describe this situation as a *liquidity trap*. According to the *IS-LM* model, expansionary monetary policy works by reducing interest rates and stimulating investment spending. But if interest rates have already fallen almost to zero, then perhaps monetary policy is no longer effective. Nominal interest rates cannot fall below zero: rather than making a loan at a negative nominal interest rate, a person would simply hold cash. In this environment, expansionary monetary policy raises the supply of money, making the public more liquid, but because interest rates can't fall any further, the extra liquidity might not have any effect. Aggregate demand, production, and employment may be “trapped” at low levels.

Other economists are skeptical about this argument. One response is that expansionary monetary policy might raise inflation expectations. Even if nominal interest rates cannot fall any further, higher expected inflation can lower real in-

terest rates by making them negative, which would stimulate investment spending. A second response is that monetary expansion would cause the currency to lose value in the market for foreign-currency exchange. This depreciation would make the nation's goods cheaper abroad, stimulating export demand. This second argument goes beyond the closed-economy *IS-LM* model we have used in this chapter, but it has merit in the open-economy version of the model developed in the next chapter.

Is the liquidity trap something about which monetary policymakers need to worry? Might the tools of monetary policy at times lose their power to influence the economy? There is no consensus about the answers. Skeptics say we shouldn't worry about the liquidity trap. But others say the possibility of a liquidity trap argues for a target rate of inflation greater than zero. Under zero inflation, the real interest rate, like the nominal interest, can never fall below zero. But if the normal rate of inflation is, say, 3 percent, then the central bank can easily push the real interest rate to negative 3 percent by lowering the nominal interest rate toward zero. Thus, moderate inflation gives monetary policymakers more room to stimulate the economy when needed, reducing the risk of falling into a liquidity trap.⁶

Although the model presented in this chapter provides the basic framework for analyzing aggregate demand, it is not the whole story. In later chapters, we examine in more detail the elements of this model and thereby refine our understanding of aggregate demand. In Chapter 16, for example, we study theories of consumption. Because the consumption function is a crucial piece of the *IS-LM* model, a deeper analysis of consumption may modify our view of the impact of monetary and fiscal policy on the economy. The simple *IS-LM* model presented in Chapters 10 and 11 provides the starting point for this further analysis.

⁶ To read more about the liquidity trap, see Paul R. Krugman, “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap,” *Brookings Panel on Economic Activity* (1998): 137–205.

Summary

1. The *IS–LM* model is a general theory of the aggregate demand for goods and services. The exogenous variables in the model are fiscal policy, monetary policy, and the price level. The model explains two endogenous variables: the interest rate and the level of national income.
2. The *IS* curve represents the negative relationship between the interest rate and the level of income that arises from equilibrium in the market for goods and services. The *LM* curve represents a positive relationship between the interest rate and the level of income that arises from equilibrium in the market for real money balances. Equilibrium in the *IS–LM* model—the intersection of the *IS* and *LM* curves—represents simultaneous equilibrium in the market for goods and services and in the market for real money balances.
3. The aggregate demand curve summarizes the results from the *IS–LM* model by showing equilibrium income at any given price level. The aggregate demand curve slopes downward because a lower price level increases real money balances, lowers the interest rate, stimulates investment spending, and thereby raises equilibrium income.
4. Expansionary fiscal policy—an increase in government purchases or a decrease in taxes—shifts the *IS* curve to the right. This shift in the *IS* curve increases the interest rate and income. The increase in income represents a rightward shift in the aggregate demand curve. Similarly, contractionary fiscal policy shifts the *IS* curve to the left, lowers the interest rate and income, and shifts the aggregate demand curve to the left.
5. Expansionary monetary policy shifts the *LM* curve downward. This shift in the *LM* curve lowers the interest rate and raises income. The increase in income represents a rightward shift of the aggregate demand curve. Similarly, contractionary monetary policy shifts the *LM* curve upward, raises the interest rate, lowers income, and shifts the aggregate demand curve to the left.

KEY CONCEPTS

Monetary transmission mechanism

Pigou effect

Debt-deflation theory

QUESTIONS FOR REVIEW

1. Explain why the aggregate demand curve slopes downward.
2. What is the impact of an increase in taxes on the interest rate, income, consumption, and investment?
3. What is the impact of a decrease in the money supply on the interest rate, income, consumption, and investment?
4. Describe the possible effects of falling prices on equilibrium income.

PROBLEMS AND APPLICATIONS

1. According to the *IS–LM* model, what happens to the interest rate, income, consumption, and investment under the following circumstances?
 - a. The central bank increases the money supply.
 - b. The government increases government purchases.
 - c. The government increases taxes.
 - d. The government increases government purchases and taxes by equal amounts.
2. Use the *IS–LM* model to predict the effects of each of the following shocks on income, the interest rate, consumption, and investment. In each case, explain what the Fed should do to keep income at its initial level.
 - a. After the invention of a new high-speed computer chip, many firms decide to upgrade their computer systems.
 - b. A wave of credit-card fraud increases the frequency with which people make transactions in cash.
 - c. A best-seller titled *Retire Rich* convinces the public to increase the percentage of their income devoted to saving.
3. Consider the economy of Hicksonia.
 - a. The consumption function is given by

$$C = 200 + 0.75(Y - T).$$

The investment function is

$$I = 200 - 25r.$$

Government purchases and taxes are both 100. For this economy, graph the *IS* curve for r ranging from 0 to 8.
 - b. The money demand function in Hicksonia is

$$(M/P)^d = Y - 100r.$$

The money supply M is 1,000 and the price level P is 2. For this economy, graph the *LM* curve for r ranging from 0 to 8.
 - c. Find the equilibrium interest rate r and the equilibrium level of income Y .
 - d. Suppose that government purchases are raised from 100 to 150. How much does the *IS* curve shift? What are the new equilibrium interest rate and level of income?
 - e. Suppose instead that the money supply is raised from 1,000 to 1,200. How much does the *LM* curve shift? What are the new equilibrium interest rate and level of income?
 - f. With the initial values for monetary and fiscal policy, suppose that the price level rises from 2 to 4. What happens? What are the new equilibrium interest rate and level of income?
 - g. Derive and graph an equation for the aggregate demand curve. What happens to this aggregate demand curve if fiscal or monetary policy changes, as in parts (d) and (e)?
4. Explain why each of the following statements is true. Discuss the impact of monetary and fiscal policy in each of these special cases.
 - a. If investment does not depend on the interest rate, the *IS* curve is vertical.
 - b. If money demand does not depend on the interest rate, the *LM* curve is vertical.
 - c. If money demand does not depend on income, the *LM* curve is horizontal.
 - d. If money demand is extremely sensitive to the interest rate, the *LM* curve is horizontal.
5. Suppose that the government wants to raise investment but keep output constant. In the *IS–LM* model, what mix of monetary and fiscal policy will achieve this goal? In the early 1980s, the U.S. government cut taxes and ran a budget deficit while the Fed pursued a tight monetary policy. What effect should this policy mix have?
6. Use the *IS–LM* diagram to describe the short-run and long-run effects of the following changes on national income, the interest rate, the price level, consumption, investment, and real money balances.
 - a. An increase in the money supply.
 - b. An increase in government purchases.
 - c. An increase in taxes.

7. The Fed is considering two alternative monetary policies:

- holding the money supply constant and letting the interest rate adjust, or
- adjusting the money supply to hold the interest rate constant.

In the *IS–LM* model, which policy will better stabilize output under the following conditions?

- a. All shocks to the economy arise from exogenous changes in the demand for goods and services.
- b. All shocks to the economy arise from exogenous changes in the demand for money.

8. Suppose that the demand for real money balances depends on disposable income. That is, the money demand function is

$$M/P = L(r, Y - T).$$

Using the *IS–LM* model, discuss whether this change in the money demand function alters the following:

- a. The analysis of changes in government purchases.
- b. The analysis of changes in taxes.

APPENDIX

The Simple Algebra of the *IS–LM* Model and the Aggregate Demand Curve

The chapter analyzes the *IS–LM* model with graphs of the *IS* and *LM* curves. Here we analyze the model algebraically rather than graphically. This alternative presentation offers additional insight into how monetary and fiscal policy influence aggregate demand.

The *IS* Curve

One way to think about the *IS* curve is that it describes the combinations of income Y and the interest rate r that satisfy an equation we first saw in Chapter 3:

$$Y = C(Y - T) + I(r) + G.$$

This equation combines the national income accounts identity, the consumption function, and the investment function. It states that the quantity of goods produced, Y , must equal the quantity of goods demanded, $C + I + G$.

We can learn more about the *IS* curve by considering the special case in which the consumption function and investment function are linear. We begin with the national income accounts identity

$$Y = C + I + G.$$

Now suppose that the consumption function is

$$C = a + b(Y - T),$$

where a and b are numbers greater than zero, and the investment function is

$$I = c - dr,$$

where c and d also are numbers greater than zero. The parameter b is the marginal propensity to consume, so we expect b to be between zero and one. The parameter d determines how much investment responds to the interest rate; because investment rises when the interest rate falls, there is a minus sign in front of d .

From these three equations, we can derive an algebraic expression for the *IS* curve and see what influences the *IS* curve's position and slope. If we substitute the consumption and investment functions into the national income accounts identity, we obtain

$$Y = [a + b(Y - T)] + (c - dr) + G.$$

Note that Y shows up on both sides of this equation. We can simplify this equation by bringing all the Y terms to the left-hand side and rearranging the terms on the right-hand side:

$$Y - bY = (a + c) + (G - bT) - dr.$$

We solve for Y to get

$$Y = \frac{a + c}{1 - b} + \frac{1}{1 - b}G + \frac{-b}{1 - b}T + \frac{-d}{1 - b}r.$$

This equation expresses the *IS* curve algebraically. It tells us the level of income Y for any given interest rate r and fiscal policy G and T . Holding fiscal policy fixed, the equation gives us a relationship between the interest rate and the level of income: the higher the interest rate, the lower the level of income. The *IS* curve graphs this equation for different values of Y and r given fixed values of G and T .

Using this last equation, we can verify our previous conclusions about the *IS* curve. First, because the coefficient of the interest rate is negative, the *IS* curve slopes downward: higher interest rates reduce income. Second, because the coefficient of government purchases is positive, an increase in government purchases shifts the *IS* curve to the right. Third, because the coefficient of taxes is negative, an increase in taxes shifts the *IS* curve to the left.

The coefficient of the interest rate, $-d/(1 - b)$, tells us what determines whether the *IS* curve is steep or flat. If investment is highly sensitive to the interest rate, then d is large, and income is highly sensitive to the interest rate as well. In this case, small changes in the interest rate lead to large changes in income: the *IS* curve is relatively flat. Conversely, if investment is not very sensitive to the interest rate, then d is small, and income is also not very sensitive to the interest rate. In this case, large changes in interest rates lead to small changes in income: the *IS* curve is relatively steep.

Similarly, the slope of the *IS* curve depends on the marginal propensity to consume b . The larger the marginal propensity to consume, the larger the change in income resulting from a given change in the interest rate. The reason is that a large marginal propensity to consume leads to a large multiplier for changes in investment. The larger the multiplier, the larger the impact of a change in investment on income, and the flatter the *IS* curve.

The marginal propensity to consume b also determines how much changes in fiscal policy shift the *IS* curve. The coefficient of G , $1/(1 - b)$, is the government-purchases multiplier in the Keynesian cross. Similarly, the coefficient of T , $-b/(1 - b)$, is the tax multiplier in the Keynesian cross. The larger the marginal propensity to consume, the greater the multiplier, and thus the greater the shift in the *IS* curve that arises from a change in fiscal policy.

The *LM* Curve

The *LM* curve describes the combinations of income Y and the interest rate r that satisfy the money market equilibrium condition

$$M/P = L(r, Y).$$

This equation simply equates money supply and money demand.

We can learn more about the LM curve by considering the case in which the money demand function is linear—that is,

$$L(r, Y) = eY - fr,$$

where e and f are numbers greater than zero. The value of e determines how much the demand for money rises when income rises. The value of f determines how much the demand for money falls when the interest rate rises. There is a minus sign in front of the interest rate term because money demand is inversely related to the interest rate.

The equilibrium in the money market is now described by

$$M/P = eY - fr.$$

To see what this equation implies, rearrange the terms so that r is on the left-hand side. We obtain

$$r = (e/f)Y - (1/f)M/P.$$

This equation gives us the interest rate that equilibrates the money market for any values of income and real money balances. The LM curve graphs this equation for different values of Y and r given a fixed value of M/P .

From this last equation, we can verify some of our conclusions about the LM curve. First, because the coefficient of income is positive, the LM curve slopes upward: higher income requires a higher interest rate to equilibrate the money market. Second, because the coefficient of real money balances is negative, decreases in real balances shift the LM curve upward, and increases in real balances shift the LM curve downward.

From the coefficient of income, e/f , we can see what determines whether the LM curve is steep or flat. If money demand is not very sensitive to the level of income, then e is small. In this case, only a small change in the interest rate is necessary to offset the small increase in money demand caused by a change in income: the LM curve is relatively flat. Similarly, if the quantity of money demanded is not very sensitive to the interest rate, then f is small. In this case, a shift in money demand caused by a change in income leads to a large change in the equilibrium interest rate: the LM curve is relatively steep.

The Aggregate Demand Curve

To find the aggregate demand equation, we must find the level of income that satisfies both the IS equation and the LM equation. To do this, substitute the LM equation for the interest rate r into the IS equation to obtain

$$Y = \frac{a + c}{1 - b} + \frac{1}{1 - b}G + \frac{-b}{1 - b}T + \frac{-d}{1 - b} \left(\frac{e}{f}Y - \frac{1}{f} \frac{M}{P} \right).$$

With some algebraic manipulation, we can solve for Y . The final equation for Y is

$$Y = \frac{z(a+c)}{1-b} + \frac{z}{1-b}G + \frac{-zb}{1-b}T + \frac{d}{(1-b)[f+de/(1-b)]} \frac{M}{P},$$

where $z = f/[f + de/(1-b)]$ is a composite of some of the parameters and is between zero and one.

This last equation expresses the aggregate demand curve algebraically. It says that income depends on fiscal policy G and T , monetary policy M , and the price level P . The aggregate demand curve graphs this equation for different values of Y and P given fixed values of G , T , and M .

We can explain the slope and position of the aggregate demand curve with this equation. First, the aggregate demand curve slopes downward, because an increase in P lowers M/P and thus lowers Y . Second, increases in the money supply raise income and shift the aggregate demand curve to the right. Third, increases in government purchases or decreases in taxes also raise income and shift the aggregate demand curve to the right. Note that, because z is less than one, the multipliers for fiscal policy are smaller in the $IS-LM$ model than in the Keynesian cross. Hence, the parameter z reflects the crowding out of investment discussed earlier.

Finally, this equation shows the relationship between the aggregate demand curve derived in this chapter from the $IS-LM$ model and the aggregate demand curve derived in Chapter 9 from the quantity theory of money. The quantity theory assumes that the interest rate does not influence the quantity of real money balances demanded. Put differently, the quantity theory assumes that the parameter f equals zero. If f equals zero, then the composite parameter z also equals zero, so fiscal policy does not influence aggregate demand. Thus, the aggregate demand curve derived in Chapter 9 is a special case of the aggregate demand curve derived here.

CASE STUDY

The Effectiveness of Monetary and Fiscal Policy

Economists have long debated whether monetary or fiscal policy exerts a more powerful influence on aggregate demand. According to the $IS-LM$ model, the answer to this question depends on the parameters of the IS and LM curves. Therefore, economists have spent much energy arguing about the size of these parameters. The most hotly contested parameters are those that describe the influence of the interest rate on economic decisions.

Those economists who believe that fiscal policy is more potent than monetary policy argue that the responsiveness of investment to the interest rate—measured by the parameter d —is small. If you look at the algebraic equation for aggregate demand, you will see that a small value of d implies a small effect of the money supply on income. The reason is that when d is small, the IS curve is nearly vertical, and shifts in the LM curve do not cause much of a

change in income. In addition, a small value of d implies a large value of z , which in turn implies that fiscal policy has a large effect on income. The reason for this large effect is that when investment is not very responsive to the interest rate, there is little crowding out.

Those economists who believe that monetary policy is more potent than fiscal policy argue that the responsiveness of money demand to the interest rate—measured by the parameter f —is small. When f is small, z is small, and fiscal policy has a small effect on income; in this case, the LM curve is nearly vertical. In addition, when f is small, changes in the money supply have a large effect on income.

Few economists today endorse either of these extreme views. The evidence indicates that the interest rate affects both investment and money demand. This finding implies that both monetary and fiscal policy are important determinants of aggregate demand.

MORE PROBLEMS AND APPLICATIONS

1. Give an algebraic answer to each of the following questions. Then explain in words the economics that underlies your answer.
 - a. How does the sensitivity of investment to the interest rate affect the slope of the aggregate demand curve?
 - b. How does the sensitivity of money demand to the interest rate affect the slope of the aggregate demand curve?
 - c. How does the marginal propensity to consume affect the response of aggregate demand to changes in government purchases?

Aggregate Demand in the Open Economy

When conducting monetary and fiscal policy, policymakers often look beyond their own country's borders. Even if domestic prosperity is their sole objective, it is necessary for them to consider the rest of the world. The international flow of goods and services and the international flow of capital can affect an economy in profound ways. Policymakers ignore these effects at their peril.

In this chapter we extend our analysis of aggregate demand to include international trade and finance. The model developed in this chapter, called the **Mundell–Fleming model**, is an open-economy version of the *IS–LM* model. Both models stress the interaction between the goods market and the money market. Both models assume that the price level is fixed and then show what causes short-run fluctuations in aggregate income (or, equivalently, shifts in the aggregate demand curve). The key difference is that the *IS–LM* model assumes a closed economy, whereas the Mundell–Fleming model assumes an open economy. The Mundell–Fleming model extends the short-run model of national income from Chapters 10 and 11 by including the effects of international trade and finance from Chapter 5.

The Mundell–Fleming model makes one important and extreme assumption: it assumes that the economy being studied is a small open economy with perfect capital mobility. That is, the economy can borrow or lend as much as it wants in world financial markets and, as a result, the economy's interest rate is determined by the world interest rate. One virtue of this assumption is that it simplifies the analysis: once the interest rate is determined, we can concentrate our attention on the role of the exchange rate. In addition, for some economies, such as Belgium or the Netherlands, the assumption of a small open economy with perfect capital mobility is a good one. Yet this assumption—and thus the Mundell–Fleming model—does not apply exactly to a large open economy such as the United States. In the conclusion to this chapter (and more fully in the appendix), we consider what happens in the more complex case in which international capital mobility is less than perfect or a nation is so large it can influence world financial markets.

One lesson from the Mundell–Fleming model is that the behavior of an economy depends on the exchange-rate system it has adopted. We begin by assuming that the economy operates with a floating exchange rate. That is, we assume that the central bank allows the exchange rate to adjust to changing economic conditions. We then examine how the economy operates under a fixed exchange rate,

and we discuss whether a floating or fixed exchange rate is better. This question has been important in recent years, as many nations around the world have debated what exchange-rate system to adopt.

12-1 The Mundell–Fleming Model

In this section we build the Mundell–Fleming model, and in the following sections we use the model to examine the impact of various policies. As you will see, the Mundell–Fleming model is built from components we have used in previous chapters. But these pieces are put together in a new way to address a new set of questions.¹

The Key Assumption: Small Open Economy With Perfect Capital Mobility

Let's begin with the assumption of a small open economy with perfect capital mobility. As we saw in Chapter 5, this assumption means that the interest rate in this economy r is determined by the world interest rate r^* . Mathematically, we can write this assumption as

$$r = r^*.$$

This world interest rate is assumed to be exogenously fixed because the economy is sufficiently small relative to the world economy that it can borrow or lend as much as it wants in world financial markets without affecting the world interest rate.

Although the idea of perfect capital mobility is expressed with a simple equation, it is important not to lose sight of the sophisticated process that this equation represents. Imagine that some event were to occur that would normally raise the interest rate (such as a decline in domestic saving). In a small open economy, the domestic interest rate might rise by a little bit for a short time, but as soon as it did, foreigners would see the higher interest rate and start lending to this country (by, for instance, buying this country's bonds). The capital inflow would drive the domestic interest rate back toward r^* . Similarly, if any event were ever to start driving the domestic interest rate downward, capital would flow out of the country to earn a higher return abroad, and this capital outflow would drive the domestic interest rate back upward toward r^* . Hence, the $r = r^*$ equation represents the assumption that the international flow of capital is rapid enough to keep the domestic interest rate equal to the world interest rate.

¹ The Mundell–Fleming model was developed in the early 1960s. Mundell's contributions are collected in Robert A. Mundell, *International Economics* (New York: Macmillan, 1968). For Fleming's contribution, see J. Marcus Fleming, "Domestic Financial Policies Under Fixed and Under Floating Exchange Rates," *IMF Staff Papers* 9 (November 1962): 369–379. In 1999, Robert Mundell was awarded the Nobel Prize for his work in open-economy macroeconomics.

The Goods Market and the IS^* Curve

The Mundell–Fleming model describes the market for goods and services much as the IS – LM model does, but it adds a new term for net exports. In particular, the goods market is represented with the following equation:

$$Y = C(Y - T) + I(r^*) + G + NX(e).$$

This equation states that aggregate income Y is the sum of consumption C , investment I , government purchases G , and net exports NX . Consumption depends positively on disposable income $Y - T$. Investment depends negatively on the interest rate, which equals the world interest rate r^* . Net exports depend negatively on the exchange rate e . As before, we define the exchange rate e as the amount of foreign currency per unit of domestic currency—for example, e might be 100 yen per dollar.

You may recall that in Chapter 5 we related net exports to the real exchange rate (the relative price of goods at home and abroad) rather than the nominal exchange rate (the relative price of domestic and foreign currencies). If e is the nominal exchange rate, then the real exchange rate ϵ equals eP/P^* , where P is the domestic price level and P^* is the foreign price level. The Mundell–Fleming model, however, assumes that the price levels at home and abroad are fixed, so the real exchange rate is proportional to the nominal exchange rate. That is, when the nominal exchange rate appreciates (say, from 100 to 120 yen per dollar), foreign goods become cheaper compared to domestic goods, and this causes exports to fall and imports to rise.

We can illustrate this equation for goods market equilibrium on a graph in which income is on the horizontal axis and the exchange rate is on the vertical axis. This curve is shown in panel (c) of Figure 12-1 and is called the IS^* curve. The new label reminds us that the curve is drawn holding the interest rate constant at the world interest rate r^* .

The IS^* curve slopes downward because a higher exchange rate reduces net exports, which in turn lowers aggregate income. To show how this works, the other panels of Figure 12-1 combine the net-exports schedule and the Keynesian cross to derive the IS^* curve. In panel (a), an increase in the exchange rate from e_1 to e_2 lowers net exports from $NX(e_1)$ to $NX(e_2)$. In panel (b), the reduction in net exports shifts the planned-expenditure schedule downward and thus lowers income from Y_1 to Y_2 . The IS^* curve summarizes this relationship between the exchange rate e and income Y .


The Money Market and the LM^* Curve

The Mundell–Fleming model represents the money market with an equation that should be familiar from the IS – LM model, with the additional assumption that the domestic interest rate equals the world interest rate:

$$M/P = L(r^*, Y).$$

figure 12-1


The IS^* Curve The IS^* curve is derived from the net-exports schedule and the Keynesian cross. Panel (a) shows the net-exports schedule: an increase in the exchange rate from e_1 to e_2 lowers net exports from $NX(e_1)$ to $NX(e_2)$. Panel (b) shows the Keynesian cross: a decrease in net exports from $NX(e_1)$ to $NX(e_2)$ shifts the planned-expenditure schedule downward and reduces income from Y_1 to Y_2 . Panel (c) shows the IS^* curve summarizing this relationship between the exchange rate and income: the higher the exchange rate, the lower the level of income.



This equation states that the supply of real money balances, M/P , equals the demand, $L(r, Y)$. The demand for real balances depends negatively on the interest rate, which is now set equal to the world interest rate r^* , and positively on income Y . The money supply M is an exogenous variable controlled by the central bank, and because the Mundell–Fleming model is designed to analyze short-run fluctuations, the price level P is also assumed to be exogenously fixed.

We can represent this equation graphically with a vertical LM^* curve, as in panel (b) of Figure 12-2. The LM^* curve is vertical because the exchange rate does not enter into the LM^* equation. Given the world interest rate, the LM^* equation determines aggregate income, regardless of the exchange rate. Figure 12-2 shows how the LM^* curve arises from the world interest rate and the LM curve, which relates the interest rate and income.

figure 12-2



The LM^* Curve Panel (a) shows the standard LM curve [which graphs the equation $M/P = L(r, Y)$] together with a horizontal line representing the world interest rate r^* . The intersection of these two curves determines the level of income, regardless of the exchange rate. Therefore, as panel (b) shows, the LM^* curve is vertical.

Putting the Pieces Together


According to the Mundell–Fleming model, a small open economy with perfect capital mobility can be described by two equations:

$$\begin{aligned} Y &= C(Y - T) + I(r^*) + G + NX(e) & IS^*, \\ M/P &= L(r^*, Y) & LM^*. \end{aligned}$$

The first equation describes equilibrium in the goods market, and the second equation describes equilibrium in the money market. The exogenous variables are fiscal policy G and T , monetary policy M , the price level P , and the world interest rate r^* . The endogenous variables are income Y and the exchange rate e .

Figure 12-3 illustrates these two relationships. The equilibrium for the economy is found where the IS^* curve and the LM^* curve intersect. This intersection

figure 12-3

**The Mundell–Fleming Model**

This graph of the Mundell–Fleming model plots the goods market equilibrium condition IS^* and the money market equilibrium condition LM^* . Both curves are drawn holding the interest rate constant at the world interest rate. The intersection of these two curves shows the level of income and the exchange rate that satisfy equilibrium both in the goods market and in the money market.

shows the exchange rate and the level of income at which both the goods market and the money market are in equilibrium. With this diagram, we can use the Mundell–Fleming model to show how aggregate income Y and the exchange rate e respond to changes in policy.

12-2 The Small Open Economy Under Floating Exchange Rates


Before analyzing the impact of policies in an open economy, we must specify the international monetary system in which the country has chosen to operate. We start with the system relevant for most major economies today: **floating exchange rates**. Under floating exchange rates, the exchange rate is allowed to fluctuate in response to changing economic conditions.

Fiscal Policy

Suppose that the government stimulates domestic spending by increasing government purchases or by cutting taxes. Because such expansionary fiscal policy increases planned expenditure, it shifts the IS^* curve to the right, as in Figure 12-4. As a result, the exchange rate appreciates, whereas the level of income remains the same.

Notice that fiscal policy has very different effects in a small open economy than it does in a closed economy. In the closed-economy $IS-LM$ model, a fiscal expansion raises income, whereas in a small open economy with a floating exchange rate, a fiscal expansion leaves income at the same level. Why the

figure 12-4



A Fiscal Expansion Under Floating Exchange Rates An increase in government purchases or a decrease in taxes shifts the IS^* curve to the right. This raises the exchange rate but has no effect on income.

difference? When income rises in a closed economy, the interest rate rises, because higher income increases the demand for money. That is not possible in a small open economy: as soon as the interest rate tries to rise above the world interest rate r^* , capital flows in from abroad. This capital inflow increases the demand for the domestic currency in the market for foreign-currency exchange and, thus, bids up the value of the domestic currency. The appreciation of the exchange rate makes domestic goods expensive relative to foreign goods, and this reduces net exports. The fall in net exports offsets the effects of the expansionary fiscal policy on income.

Why is the fall in net exports so great that it renders fiscal policy powerless to influence income? To answer this question, consider the equation that describes the money market:


$$M/P = L(r, Y).$$

In both closed and open economies, the quantity of real money balances supplied M/P is fixed, and the quantity demanded (determined by r and Y) must equal this fixed supply. In a closed economy, a fiscal expansion causes the equilibrium interest rate to rise. This increase in the interest rate (which reduces the quantity of money demanded) allows equilibrium income to rise (which increases the quantity of money demanded). By contrast, in a small open economy, r is fixed at r^* , so there is only one level of income that can satisfy this equation, and this level of income does not change when fiscal policy changes. Thus, when the government increases spending or cuts taxes, the appreciation of the exchange rate and the fall in net exports must be large enough to offset fully the normal expansionary effect of the policy on income.

Monetary Policy

Suppose now that the central bank increases the money supply. Because the price level is assumed to be fixed, the increase in the money supply means an increase in real balances. The increase in real balances shifts the LM^* curve to the right, as in Figure 12-5. Hence, an increase in the money supply raises income and lowers the exchange rate.

figure 12-5



A Monetary Expansion Under Floating Exchange Rates An increase in the money supply shifts the LM^* curve to the right, lowering the exchange rate and raising income.


Although monetary policy influences income in an open economy, as it does in a closed economy, the monetary transmission mechanism is different. Recall that in a closed economy an increase in the money supply increases spending because it lowers the interest rate and stimulates investment. In a small open economy, the interest rate is fixed by the world interest rate. As soon as an increase in the money supply puts downward pressure on the domestic interest rate, capital flows out of the economy as investors seek a higher return elsewhere. This capital outflow prevents the domestic interest rate from falling. In addition, because the capital outflow increases the supply of the domestic currency in the market for foreign-currency exchange, the exchange rate depreciates. The fall in the exchange rate makes domestic goods inexpensive relative to foreign goods and, thereby, stimulates net exports. Hence, in a small open economy, monetary policy influences income by altering the exchange rate rather than the interest rate.

Trade Policy

Suppose that the government reduces the demand for imported goods by imposing an import quota or a tariff. What happens to aggregate income and the exchange rate?

Because net exports equal exports minus imports, a reduction in imports means an increase in net exports. That is, the net-exports schedule shifts to the

figure 12-6



A Trade Restriction Under Floating Exchange Rates A tariff or an import quota shifts the net-exports schedule in panel (a) to the right. As a result, the IS^* curve in panel (b) shifts to the right, raising the exchange rate and leaving income unchanged.

right, as in Figure 12-6. This shift in the net-exports schedule increases planned expenditure and thus moves the IS^* curve to the right. Because the LM^* curve is vertical, the trade restriction raises the exchange rate but does not affect income.

Often a stated goal of policies to restrict trade is to alter the trade balance NX . Yet, as we first saw in Chapter 5, such policies do not necessarily have that effect. The same conclusion holds in the Mundell–Fleming model under floating exchange rates. Recall that

$$NX(e) = Y - C(Y - T) - I(r^*) - G.$$

Because a trade restriction does not affect income, consumption, investment, or government purchases, it does not affect the trade balance. Although the shift in the net-exports schedule tends to raise NX , the increase in the exchange rate reduces NX by the same amount.

12-3 The Small Open Economy Under Fixed Exchange Rates

We now turn to the second type of exchange-rate system: **fixed exchange rates**. In the 1950s and 1960s, most of the world's major economies, including the United States, operated within the Bretton Woods system—an international monetary system under which most governments agreed to fix exchange rates. The world abandoned this system in the early 1970s, and exchange rates were allowed to float. Some European countries later reinstated a system of fixed exchange rates among themselves, and some economists have advocated a return to a worldwide system of fixed exchange rates. In this section we discuss how such a system works, and we examine the impact of economic policies on an economy with a fixed exchange rate.


How a Fixed-Exchange-Rate System Works

Under a system of fixed exchange rates, a central bank stands ready to buy or sell the domestic currency for foreign currencies at a predetermined price. For example, suppose that the Fed announced that it was going to fix the exchange rate at 100 yen per dollar. It would then stand ready to give \$1 in exchange for 100 yen or to give 100 yen in exchange for \$1. To carry out this policy, the Fed would need a reserve of dollars (which it can print) and a reserve of yen (which it must have purchased previously).

A fixed exchange rate dedicates a country's monetary policy to the single goal of keeping the exchange rate at the announced level. In other words, the essence of a fixed-exchange-rate system is the commitment of the central bank to allow the money supply to adjust to whatever level will ensure that the equilibrium exchange rate equals the announced exchange rate. Moreover, as long as the central bank stands ready to buy or sell foreign currency at the fixed exchange rate, the money supply adjusts automatically to the necessary level.

To see how fixing the exchange rate determines the money supply, consider the following example. Suppose that the Fed announces that it will fix the exchange rate at 100 yen per dollar, but, in the current equilibrium with the current money supply, the exchange rate is 150 yen per dollar. This situation is illustrated in panel (a) of Figure 12-7. Notice that there is a profit opportunity: an arbitrageur could buy 300 yen in the marketplace for \$2, and then sell the yen to the Fed for \$3, making a \$1 profit. When the Fed buys these yen from the arbitrageur, the dollars it pays for them automatically increase the money supply. The

figure 12-7



How a Fixed Exchange Rate Governs the Money Supply In panel (a), the equilibrium exchange rate initially exceeds the fixed level. Arbitrageurs will buy foreign currency in foreign-exchange markets and sell it to the Fed for a profit. This process automatically increases the money supply, shifting the LM^* curve to the right and lowering the exchange rate. In panel (b), the equilibrium exchange rate is initially below the fixed level. Arbitrageurs will buy dollars in foreign-exchange markets and use them to buy foreign currency from the Fed. This process automatically reduces the money supply, shifting the LM^* curve to the left and raising the exchange rate.

rise in the money supply shifts the LM^* curve to the right, lowering the equilibrium exchange rate. In this way, the money supply continues to rise until the equilibrium exchange rate falls to the announced level.

Conversely, suppose that when the Fed announces that it will fix the exchange rate at 100 yen per dollar, the equilibrium is 50 yen per dollar. Panel (b) of Figure 12-7 shows this situation. In this case, an arbitrageur could make a profit by buying 100 yen from the Fed for \$1 and then selling the yen in the marketplace for \$2. When the Fed sells these yen, the \$1 it receives automatically reduces the money supply. The fall in the money supply shifts the LM^* curve to the left, raising the equilibrium exchange rate. The money supply continues to fall until the equilibrium exchange rate rises to the announced level.

It is important to understand that this exchange-rate system fixes the nominal exchange rate. Whether it also fixes the real exchange rate depends on the time horizon under consideration. If prices are flexible, as they are in the long run, then the real exchange rate can change even while the nominal exchange rate is fixed. Therefore, in the long run described in Chapter 5, a policy to fix the nominal exchange rate would not influence any real variable, including the real exchange rate. A fixed nominal exchange rate would influence only the money supply and the price level. Yet in the short run described by the Mundell–Fleming model, prices are fixed, so a fixed nominal exchange rate implies a fixed real exchange rate as well.

CASE STUDY**The International Gold Standard**

During the late nineteenth and early twentieth centuries, most of the world's major economies operated under a gold standard. Each country maintained a reserve of gold and agreed to exchange one unit of its currency for a specified amount of gold. Through the gold standard, the world's economies maintained a system of fixed exchange rates.

To see how an international gold standard fixes exchange rates, suppose that the U.S. Treasury stands ready to buy or sell 1 ounce of gold for \$100, and the Bank of England stands ready to buy or sell 1 ounce of gold for 100 pounds. Together, these policies fix the rate of exchange between dollars and pounds: \$1 must trade for 1 pound. Otherwise, the law of one price would be violated, and it would be profitable to buy gold in one country and sell it in the other.

For example, suppose that the exchange rate were 2 pounds per dollar. In this case, an arbitrageur could buy 200 pounds for \$100, use the pounds to buy 2 ounces of gold from the Bank of England, bring the gold to the United States, and sell it to the Treasury for \$200—making a \$100 profit. Moreover, by bringing the gold to the United States from England, the arbitrageur would increase the money supply in the United States and decrease the money supply in England.


Thus, during the era of the gold standard, the international transport of gold by arbitrageurs was an automatic mechanism adjusting the money supply and stabilizing exchange rates. This system did not completely fix exchange rates, because shipping gold across the Atlantic was costly. Yet the international gold standard did keep the exchange rate within a range dictated by transportation costs. It thereby prevented large and persistent movements in exchange rates.²

Fiscal Policy

Let's now examine how economic policies affect a small open economy with a fixed exchange rate. Suppose that the government stimulates domestic spending by increasing government purchases or by cutting taxes. This policy shifts the IS^* curve to the right, as in Figure 12-8, putting upward pressure on the exchange rate. But because the central bank stands ready to trade foreign and domestic currency at the fixed exchange rate, arbitrageurs quickly respond to the rising exchange rate by selling foreign currency to the central bank, leading to an automatic monetary expansion. The rise in the money supply shifts the LM^* curve to the right. Thus, under a fixed exchange rate, a fiscal expansion raises aggregate income.

² For more on how the gold standard worked, see the essays in Barry Eichengreen, ed., *The Gold Standard in Theory and History* (New York: Methuen, 1985).

figure 12-8




A Fiscal Expansion Under Fixed Exchange Rates A fiscal expansion shifts the IS^* curve to the right. To maintain the fixed exchange rate, the Fed must increase the money supply, thereby shifting the LM^* curve to the right. Hence, in contrast to the case of floating exchange rates, under fixed exchange rates a fiscal expansion raises income.

Monetary Policy

Imagine that a central bank operating with a fixed exchange rate were to try to increase the money supply—for example, by buying bonds from the public. What would happen? The initial impact of this policy is to shift the LM^* curve to the right, lowering the exchange rate, as in Figure 12-9. But, because the central bank is committed to trading foreign and domestic currency at a fixed exchange

figure 12-9



A Monetary Expansion Under Fixed Exchange Rates If the Fed tries to increase the money supply—for example, by buying bonds from the public—it will put downward pressure on the exchange rate. To maintain the fixed exchange rate, the money supply and the LM^* curve must return to their initial positions. Hence, under fixed exchange rates, normal monetary policy is ineffectual.

rate, arbitrageurs quickly respond to the falling exchange rate by selling the domestic currency to the central bank, causing the money supply and the LM^* curve to return to their initial positions. Hence, monetary policy as usually conducted is ineffectual under a fixed exchange rate. By agreeing to fix the exchange rate, the central bank gives up its control over the money supply.

A country with a fixed exchange rate can, however, conduct a type of monetary policy: it can decide to change the level at which the exchange rate is fixed. A reduction in the value of the currency is called a **devaluation**, and an increase in its value is called a **revaluation**. In the Mundell–Fleming model, a devaluation shifts the LM^* curve to the right; it acts like an increase in the money supply under a floating exchange rate. A devaluation thus expands net exports and raises aggregate income. Conversely, a revaluation shifts the LM^* curve to the left, reduces net exports, and lowers aggregate income.

CASE STUDY

Devaluation and the Recovery From the Great Depression

The Great Depression of the 1930s was a global problem. Although events in the United States may have precipitated the downturn, all of the world's major economies experienced huge declines in production and employment. Yet not all governments responded to this calamity in the same way.

One key difference among governments was how committed they were to the fixed exchange rate set by the international gold standard. Some countries, such as France, Germany, Italy, and the Netherlands, maintained the old rate of exchange between gold and currency. Other countries, such as Denmark, Finland, Norway, Sweden, and the United Kingdom, reduced the amount of gold they would pay for each unit of currency by about 50 percent. By reducing the gold content of their currencies, these governments devalued their currencies relative to those of other countries.


The subsequent experience of these two groups of countries conforms to the prediction of the Mundell–Fleming model. Those countries that pursued a policy of devaluation recovered quickly from the Depression. The lower value of the currency raised the money supply, stimulated exports, and expanded production. By contrast, those countries that maintained the old exchange rate suffered longer with a depressed level of economic activity.³

Trade Policy

Suppose that the government reduces imports by imposing an import quota or a tariff. This policy shifts the net-exports schedule to the right and thus shifts the IS^* curve to the right, as in Figure 12-10. The shift in the IS^* curve tends to

³ Barry Eichengreen and Jeffrey Sachs, "Exchange Rates and Economic Recovery in the 1930s," *Journal of Economic History* 45 (December 1985): 925–946.

figure 12-10



A Trade Restriction Under Fixed Exchange Rates A tariff or an import quota shifts the IS^* curve to the right. This induces an increase in the money supply to maintain the fixed exchange rate. Hence, aggregate income increases.

raise the exchange rate. To keep the exchange rate at the fixed level, the money supply must rise, shifting the LM^* curve to the right.

The result of a trade restriction under a fixed exchange rate is very different from that under a floating exchange rate. In both cases, a trade restriction shifts the net-exports schedule to the right, but only under a fixed exchange rate does a trade restriction increase net exports NX . The reason is that a trade restriction under a fixed exchange rate induces monetary expansion rather than an appreciation of the exchange rate. The monetary expansion, in turn, raises aggregate income. Recall the accounting identity

$$NX = S - I.$$

When income rises, saving also rises, and this implies an increase in net exports.

Policy in the Mundell–Fleming Model: A Summary

The Mundell–Fleming model shows that the effect of almost any economic policy on a small open economy depends on whether the exchange rate is floating or fixed. Table 12-1 summarizes our analysis of the short-run effects of fiscal, monetary, and trade policies on income, the exchange rate, and the trade balance. What is most striking is that all of the results are different under floating and fixed exchange rates.

To be more specific, the Mundell–Fleming model shows that the power of monetary and fiscal policy to influence aggregate income depends on the exchange-rate regime. Under floating exchange rates, only monetary policy can affect income. The usual expansionary impact of fiscal policy is offset by a rise in

table 12-1

The Mundell–Fleming Model: Summary of Policy Effects

Policy	EXCHANGE-RATE REGIME					
	FLOATING			FIXED		
	IMPACT ON:					
	Y	e	NX	Y	e	NX
Fiscal expansion	0	↑	↓	↑	0	0
Monetary expansion	↑	↓	↑	0	0	0
Import restriction	0	↑	0	↑	0	↑

Note: This table shows the direction of impact of various economic policies on income Y , the exchange rate e , and the trade balance NX . A “↑” indicates that the variable increases; a “↓” indicates that it decreases; a “0” indicates no effect. Remember that the exchange rate is defined as the amount of foreign currency per unit of domestic currency (for example, 100 yen per dollar).

the value of the currency. Under fixed exchange rates, only fiscal policy can affect income. The normal potency of monetary policy is lost because the money supply is dedicated to maintaining the exchange rate at the announced level.

12-4 Interest-Rate Differentials

So far, our analysis has assumed that the interest rate in a small open economy is equal to the world interest rate: $r = r^*$. To some extent, however, interest rates differ around the world. We now extend our analysis by considering the causes and effects of international interest-rate differentials.

Country Risk and Exchange-Rate Expectations

When we assumed earlier that the interest rate in our small open economy is determined by the world interest rate, we were applying the law of one price. We reasoned that if the domestic interest rate were above the world interest rate, people from abroad would lend to that country, driving the domestic interest rate down. And if the domestic interest rate were below the world interest rate, domestic residents would lend abroad to earn a higher return, driving the domestic interest rate up. In the end, the domestic interest rate would equal the world interest rate.

Why doesn't this logic always apply? There are two reasons.

One reason is country risk. When investors buy U.S. government bonds or make loans to U.S. corporations, they are fairly confident that they will be repaid with interest. By contrast, in some less developed countries, it is plausible to fear that a revolution or other political upheaval might lead to a default on loan repayments. Borrowers in such countries often have to pay higher interest rates to compensate lenders for this risk.

Another reason interest rates differ across countries is expected changes in the exchange rate. For example, suppose that people expect the French franc to fall in value relative to the U.S. dollar. Then loans made in francs will be repaid in a less valuable currency than loans made in dollars. To compensate for this expected fall in the French currency, the interest rate in France will be higher than the interest rate in the United States.

Thus, because of both country risk and expectations of future exchange-rate changes, the interest rate of a small open economy can differ from interest rates in other economies around the world. Let's now see how this fact affects our analysis.

Differentials in the Mundell–Fleming Model

To incorporate interest-rate differentials into the Mundell–Fleming model, we assume that the interest rate in our small open economy is determined by the world interest rate plus a risk premium θ :

$$r = r^* + \theta.$$

The risk premium is determined by the perceived political risk of making loans in a country and the expected change in the real exchange rate. For our purposes here, we can take the risk premium as exogenous in order to examine how changes in the risk premium affect the economy.

The model is largely the same as before. The two equations are


$$\begin{aligned} Y &= C(Y - T) + I(r^* + \theta) + G + NX(e) & IS^*, \\ M/P &= L(r^* + \theta, Y) & LM^*. \end{aligned}$$

For any given fiscal policy, monetary policy, price level, and risk premium, these two equations determine the level of income and exchange rate that equilibrate the goods market and the money market. Holding constant the risk premium, the tools of monetary, fiscal, and trade policy work as we have already seen.

Now suppose that political turmoil causes the country's risk premium θ to rise. The most direct effect is that the domestic interest rate r rises. The higher interest rate, in turn, has two effects. First, the IS^* curve shifts to the left, because the higher interest rate reduces investment. Second, the LM^* curve shifts to the right, because the higher interest rate reduces the demand for money, and this allows a higher level of income for any given money supply. [Recall that Y must satisfy the equation $M/P = L(r^* + \theta, Y)$.] As Figure 12-11 shows, these two shifts cause income to rise and the currency to depreciate.

This analysis has an important implication: expectations of the exchange rate are partially self-fulfilling. For example, suppose that people come to believe that the French franc will not be valuable in the future. Investors will place a larger risk premium on French assets: θ will rise in France. This expectation will drive

figure 12-11

**An Increase in the Risk Premium**

An increase in the risk premium associated with a country drives up its interest rate. Because the higher interest rate reduces investment, the IS^* curve shifts to the left. Because it also reduces money demand, the LM^* curve shifts to the right. Income rises, and the exchange rate depreciates.

up French interest rates and, as we have just seen, will drive down the value of the French currency. *Thus, the expectation that a currency will lose value in the future causes it to lose value today.*

One surprising—and perhaps inaccurate—prediction of this analysis is that an increase in country risk as measured by θ will cause the economy's income to increase. This occurs in Figure 12-11 because of the rightward shift in the LM^* curve. Although higher interest rates depress investment, the depreciation of the currency stimulates net exports by an even greater amount. As a result, aggregate income rises.

There are three reasons why, in practice, such a boom in income does not occur. First, the central bank might want to avoid the large depreciation of the domestic currency and, therefore, may respond by decreasing the money supply M . Second, the depreciation of the domestic currency may suddenly increase the price of imported goods, causing an increase in the price level P . Third, when some event increases the country risk premium θ , residents of the country might respond to the same event by increasing their demand for money (for any given income and interest rate), because money is often the safest asset available. All three of these changes would tend to shift the LM^* curve toward the left, which mitigates the fall in the exchange rate but also tends to depress income.

Thus, increases in country risk are not desirable. In the short run, they typically lead to a depreciating currency and, through the three channels just described, falling aggregate income. In addition, because a higher interest rate reduces investment, the long-run implication is reduced capital accumulation and lower economic growth.

CASE STUDY

International Financial Crisis: Mexico 1994–1995

In August 1994, a Mexican peso was worth 30 cents. A year later, it was worth only 16 cents. What explains this massive fall in the value of the Mexican currency? Country risk is a large part of the story.

At the beginning of 1994, Mexico was a country on the rise. The recent passage of the North American Free Trade Agreement (NAFTA), which reduced trade barriers among the United States, Canada, and Mexico, made many confident about the future of the Mexican economy. Investors around the world were eager to make loans to the Mexican government and to Mexican corporations.

Political developments soon changed that perception. A violent uprising in the Chiapas region of Mexico made the political situation in Mexico seem precarious. Then Luis Donaldo Colosio, the leading presidential candidate, was assassinated. The political future looked less certain, and many investors started placing a larger risk premium on Mexican assets.

At first, the rising risk premium did not affect the value of the peso, because Mexico was operating with a fixed exchange rate. As we have seen, under a fixed exchange rate, the central bank agrees to trade the domestic currency (pesos) for a foreign currency (dollars) at a predetermined rate. Thus, when an increase in the country risk premium put downward pressure on the value of the peso, the Mexican central bank had to accept pesos and pay out dollars. This automatic exchange-market intervention contracted the Mexican money supply (shifting the LM^* curve to the left) when the currency might otherwise have depreciated.

Yet Mexico's reserves of foreign currency were too small to maintain its fixed exchange rate. When Mexico ran out of dollars at the end of 1994, the Mexican government announced a devaluation of the peso. This decision had repercussions, however, because the government had repeatedly promised that it would not devalue. Investors became even more distrustful of Mexican policymakers and feared further Mexican devaluations.

Investors around the world (including those in Mexico) avoided buying Mexican assets. The country risk premium rose once again, adding to the upward pressure on interest rates and the downward pressure on the peso. The Mexican stock market plummeted. When the Mexican government needed to roll over some of its debt that was coming due, investors were unwilling to buy the new debt. Default appeared to be the government's only option. In just a few months, Mexico had gone from being a promising emerging economy to being a risky economy with a government on the verge of bankruptcy.

Then the United States stepped in. The U.S. government had three motives: to help its neighbor to the south, to prevent the massive illegal immigration that might follow government default and economic collapse, and to prevent the investor pessimism regarding Mexico from spreading to other developing countries. The U.S. government, together with the International Monetary Fund (IMF), led an international effort to bail out the Mexican government. In particular, the United States provided loan guarantees for Mexican government debt, which allowed the Mexican government to refinance the debt that was coming

due. These loan guarantees helped restore confidence in the Mexican economy, thereby reducing to some extent the country risk premium.

Although the U.S. loan guarantees may well have stopped a bad situation from getting worse, they did not prevent the Mexican meltdown of 1994–1995 from being a painful experience for the Mexican people. Not only did the Mexican currency lose much of its value, but Mexico also went through a deep recession. Fortunately, by the late 1990s, aggregate income was growing again, and the worst appeared to be over. But the lesson from this experience is clear and could well apply again in the future: changes in perceived country risk, often attributable to political instability, are an important determinant of interest rates and exchange rates in small open economies.

CASE STUDY

International Financial Crisis: Asia 1997–1998

In 1997, as the Mexican economy was recovering from its financial crisis, a similar story started to unfold in several Asian economies, including Thailand, South Korea, and especially Indonesia. The symptoms were familiar: high interest rates, falling asset values, and a depreciating currency. In Indonesia, for instance, short-term nominal interest rates rose above 50 percent, the stock market lost about 90 percent of its value (measured in U.S. dollars), and the rupiah fell against the dollar by more than 80 percent. The crisis led to rising inflation in these countries (because the depreciating currency made imports more expensive) and to falling GDP (because high interest rates and reduced confidence depressed spending). Real GDP in Indonesia fell about 13 percent in 1998, making the downturn larger than any U.S. recession since the Great Depression of the 1930s.

What sparked this firestorm? The problem began in the Asian banking systems. For many years, the governments in the Asian nations had been more involved in managing the allocation of resources—in particular, financial resources—than is true in the United States and other developed countries. Some commentators had applauded this “partnership” between government and private enterprise and had even suggested that the United States should follow the example. Over time, however, it became clear that many Asian banks had been extending loans to those with the most political clout rather than to those with the most profitable investment projects. Once rising default rates started to expose this “crony capitalism,” as it was then called, international investors started to lose confidence in the future of these economies. The risk premiums for Asian assets rose, causing interest rates to skyrocket and currencies to collapse.

International crises of confidence often involve a vicious circle that can amplify the problem. Here is one story about what happened in Asia:

1. Problems in the banking system eroded international confidence in these economies.
2. Loss of confidence raised risk premiums and interest rates.

3. Rising interest rates, together with the loss of confidence, depressed the prices of stock and other assets.
4. Falling asset prices reduced the value of collateral being used for bank loans.
5. Reduced collateral increased default rates on bank loans.
6. Greater defaults exacerbated problems in the banking system. Now return to step 1 to complete and continue the circle.

Some economists have used this vicious-circle argument to suggest that the Asian crisis was a self-fulfilling prophecy: bad things happened merely because people expected bad things to happen. Most economists, however, thought the political corruption of the banking system was a real problem, which was then compounded by this vicious circle of reduced confidence.

As the Asian crisis developed, the IMF and the United States tried to restore confidence, much as they had with Mexico a few years earlier. In particular, the IMF made loans to the Asian countries to help them over the crisis; in exchange for these loans, it exacted promises that the governments would reform their banking systems and eliminate crony capitalism. The IMF's hope was that the short-term loans and longer-term reforms would restore confidence, lower the risk premium, and turn the vicious circle into a virtuous circle. This policy seems to have worked: the Asian economies recovered quickly from their crisis.

12-5 Should Exchange Rates Be Floating or Fixed?

Having analyzed how an economy works under floating and fixed exchange rates, let's consider which exchange-rate regime is better.

Pros and Cons of Different Exchange-Rate Systems

The primary argument for a floating exchange rate is that it allows monetary policy to be used for other purposes. Under fixed rates, monetary policy is committed to the single goal of maintaining the exchange rate at its announced level. Yet the exchange rate is only one of many macroeconomic variables that monetary policy can influence. A system of floating exchange rates leaves monetary policy-makers free to pursue other goals, such as stabilizing employment or prices.

Advocates of fixed exchange rates argue that exchange-rate uncertainty makes international trade more difficult. After the world abandoned the Bretton Woods system of fixed exchange rates in the early 1970s, both real and nominal exchange rates became (and remained) much more volatile than anyone had expected. Some economists attribute this volatility to irrational and destabilizing speculation by international investors. Business executives often claim that this volatility is harmful because it increases the uncertainty that accompanies international business transactions. Yet, despite this exchange-rate volatility, the amount of world trade has continued to rise under floating exchange rates.

Advocates of fixed exchange rates sometimes argue that a commitment to a fixed exchange rate is one way to discipline a nation's monetary authority and prevent excessive growth in the money supply. Yet there are many other policy rules to which the central bank could be committed. In Chapter 14, for instance, we discuss policy rules such as targets for nominal GDP or the inflation rate. Fixing the exchange rate has the advantage of being simpler to implement than these other policy rules, because the money supply adjusts automatically, but this policy may lead to greater volatility in income and employment.

In the end, the choice between floating and fixed rates is not as stark as it may seem at first. During periods of fixed exchange rates, countries can change the value of their currency if maintaining the exchange rate conflicts too severely with other goals. During periods of floating exchange rates, countries often use formal or informal targets for the exchange rate when deciding whether to expand or contract the money supply. We rarely observe exchange rates that are completely fixed or completely floating. Instead, under both systems, stability of the exchange rate is usually one among many of the central bank's objectives.



"Then it's agreed. Until the dollar firms up, we let the clamshell float."

© The New Yorker collection 1971 Ed Fisher from cartoonbank.com. All Rights Reserved.

CASE STUDY

Monetary Union in the United States and Europe

If you have ever driven the 3,000 miles from New York City to San Francisco, you may recall that you never needed to change your money from one form of currency to another. In all fifty U.S. states, local residents are happy to accept the U.S. dollar for the items you buy. Such a *monetary union* is the most extreme form of a fixed exchange rate. The exchange rate between New York dollars and San Francisco dollars is so irrevocably fixed that you may not even know that there is a difference between the two. (What's the difference? Each dollar bill is issued by one of the dozen local Federal Reserve Banks. Although the bank of origin can be identified from the bill's markings, you don't care which type of dollar you hold because everyone else, including the Federal Reserve system, is ready to trade them one for one.)

If you have ever made a similar 3,000-mile trip across Europe, however, your experience was probably very different. You didn't have to travel far before needing to exchange your French francs for German marks, Dutch guilders, Spanish pesetas, or Italian lira. The large number of currencies in Europe made traveling less convenient and more expensive. Every time you crossed a border, you had to wait in line at a bank to get the local money, and you had to pay the bank a fee for the service.

Recently, however, this has started to change. Many countries in Europe have decided to form their own monetary union and use a common currency called

the *euro*, which was introduced in January 1999. The adoption of the euro is an extension of the *European Monetary System (EMS)*, which during the previous two decades had attempted to limit exchange-rate fluctuations among participating countries. When the euro is fully adopted, this goal will be achieved: the exchange rate between France and Germany will be as fixed as the exchange rate between New York City and San Francisco.

The introduction of a common currency has its costs. The most important is that the nations of Europe will no longer be able to conduct their own monetary policies. Instead, a European central bank, with participation of all member countries, will set a single monetary policy for all of Europe. The central banks of the individual countries will play a role similar to that of regional Federal Reserve Banks: they will monitor local conditions but they will have no control over the money supply or interest rates. Critics of the move toward a common currency argue that the cost of losing national monetary policy is large. If a recession hits one country but not others in Europe, that country may wish it had the tool of monetary policy to combat the downturn.

Why, according to these economists, is monetary union a bad idea for Europe if it works so well in the United States? These economists argue that the United States is different from Europe in two important ways. First, labor is more mobile among U.S. states than among European countries. This is in part because the United States has a common language and in part because most Americans are descended from immigrants, who have shown a willingness to move. Therefore, when a regional recession occurs, U.S. workers are more likely to move from high-unemployment states to low-unemployment states. Second, the United States has a strong central government that can use fiscal policy—such as the federal income tax—to redistribute resources among regions. Because Europe does not have these two advantages, it will suffer more when it restricts itself to a single monetary policy.

Advocates of a common currency believe that the loss of national monetary policy is more than offset by other gains. With a single currency in all of Europe, travelers and businesses will no longer need to worry about exchange rates, and this should encourage more international trade. In addition, a common currency may have the political advantage of making Europeans feel more connected to one another. The twentieth century was marked by two world wars, both of which were sparked by European discord. If a common currency makes the nations of Europe more harmonious, it will benefit the entire world.

Speculative Attacks, Currency Boards, and Dollarization

Imagine that you are a central banker of a small country. You and your fellow policymakers decide to fix your currency—let's call it the peso—against the U.S. dollar. From now on, one peso will sell for one dollar.

As we discussed earlier, you now have to stand ready to buy and sell pesos for a dollar each. The money supply will adjust automatically to make the equilibrium exchange rate equal your target. There is, however, one potential problem with this plan: you might run out of dollars. If people come to the central bank

to sell large quantities of pesos, the central bank's dollar reserves might dwindle to zero. In this case, the central bank has no choice but to abandon the fixed exchange rate and let the peso depreciate.

This fact raises the possibility of a *speculative attack*—a change in investors' perceptions that makes the fixed exchange rate untenable. Suppose that, for no good reason, a rumor spreads that the central bank is going to abandon the exchange-rate peg. People would respond by rushing to the central bank to convert pesos into dollars before the pesos lose value. This rush would drain the central bank's reserves and could force the central bank to abandon the peg. In this case, the rumor would prove self-fulfilling.

To avoid this possibility, some economists argue that a fixed exchange rate should be supported by a *currency board*, such as that used by Argentina in the 1990s. A currency board is an arrangement by which the central bank holds enough foreign currency to back each unit of the domestic currency. In our example, the central bank would hold one U.S. dollar (or one dollar invested in a U.S. government bond) for every peso. No matter how many pesos turned up at the central bank to be exchanged, the central bank would never run out of dollars.

Once a central bank has adopted a currency board, it might consider the natural next step: it can abandon the peso altogether and let its country use the U.S. dollar. Such a plan is called *dollarization*. It happens on its own in high-inflation economies, where foreign currencies offer a more reliable store of value than the domestic currency. But it can also occur as a matter of public policy: Panama is an example. If a country really wants its currency to be irrevocably fixed to the dollar, the most reliable method is to make its currency the dollar. The only loss from dollarization is the small seigniorage revenue, which accrues to the U.S. government.⁴

12-6 The Mundell–Fleming Model With a Changing Price Level

So far we have been using the Mundell–Fleming model to study the small open economy in the short run when the price level is fixed. To see how this model relates to models we have examined previously, let's consider what happens when the price level changes.

To examine price adjustment in an open economy, we must distinguish between the nominal exchange rate e and the real exchange rate ϵ , which equals eP/P^* . We can write the Mundell–Fleming model as

$$\begin{aligned} Y &= C(Y - T) + I(r^*) + G + NX(\epsilon) & IS^*, \\ M/P &= L(r^*, Y) & LM^*. \end{aligned}$$


⁴ Dollarization may also lead to a loss in national pride from seeing American portraits on the currency. If it wanted, the U.S. government could fix this problem by leaving blank the center space that now has George Washington's portrait. Each nation using the U.S. dollar could insert the face of its own local hero.

These equations should be familiar by now. The first equation describes the IS^* curve, and the second equation describes the LM^* curve. Note that net exports depend on the real exchange rate.

Figure 12-12 shows what happens when the price level falls. Because a lower price level raises the level of real money balances, the LM^* curve shifts to the right, as in panel (a) of Figure 12-12. The real exchange rate depreciates, and the equilibrium level of income rises. The aggregate demand curve summarizes this negative relationship between the price level and the level of income, as shown in panel (b) of Figure 12-12.

Thus, just as the $IS-LM$ model explains the aggregate demand curve in a closed economy, the Mundell-Fleming model explains the aggregate demand curve for a small open economy. In both cases, the aggregate demand curve shows the set of equilibria that arise as the price level varies. And in both cases, anything that changes the equilibrium for a given price level shifts the aggregate demand curve. Policies that raise income shift the aggregate demand curve to the right; policies that lower income shift the aggregate demand curve to the left.


figure 12-12



Mundell-Fleming as a Theory of Aggregate Demand

Panel (a) shows that when the price level falls, the LM^* curve shifts to the right. The equilibrium level of income rises. Panel (b) shows that this negative relationship between P and Y is summarized by the aggregate demand curve.

figure 12-13



The Short-Run and Long-Run Equilibria in a Small Open Economy

Point K in both panels shows the equilibrium under the Keynesian assumption that the price level is fixed at P_1 . Point C in both panels shows the equilibrium under the classical assumption that the price level adjusts to maintain income at its natural rate $\bar{Y}$.

We can use this diagram to show how the short-run model in this chapter is related to the long-run model in Chapter 5. Figure 12-13 shows the short-run and long-run equilibria. In both panels of the figure, point K describes the short-run equilibrium, because it assumes a fixed price level. At this equilibrium, the demand for goods and services is too low to keep the economy producing at its natural rate. Over time, low demand causes the price level to fall. The fall in the price level raises real money balances, shifting the LM^* curve to the right. The real exchange rate depreciates, so net exports rise. Eventually, the economy reaches point C, the long-run equilibrium. The speed of transition between the short-run and long-run equilibria depends on how quickly the price level adjusts to restore the economy to the natural rate.

The levels of income at point K and point C are both of interest. Our central concern in this chapter has been how policy influences point K, the short-run equilibrium. In Chapter 5 we examined the determinants of point C, the long-run equilibrium. Whenever policymakers consider any change in policy, they need to consider both the short-run and long-run effects of their decision.

12-7 A Concluding Reminder

In this chapter we have examined how a small open economy works in the short run when prices are sticky. We have seen how monetary and fiscal policy influence income and the exchange rate, and how the behavior of the economy depends on whether the exchange rate is floating or fixed. In closing, it is worth repeating a lesson from Chapter 5. Many countries, including the United States, are neither closed economies nor small open economies: they lie somewhere in between.

A large open economy, such as the United States, combines the behavior of a closed economy and the behavior of a small open economy. When analyzing policies in a large open economy, we need to consider both the closed-economy logic of Chapter 11 and the open-economy logic developed in this chapter. The appendix to this chapter presents a model for a large open economy. The results of that model are, as one would guess, a mixture of the two polar cases we have already examined.

To see how we can draw on the logic of both the closed and small open economies and apply these insights to the United States, consider how a monetary contraction affects the economy in the short run. In a closed economy, a monetary contraction raises the interest rate, lowers investment, and thus lowers aggregate income. In a small open economy with a floating exchange rate, a monetary contraction raises the exchange rate, lowers net exports, and thus lowers aggregate income. The interest rate is unaffected, however, because it is determined by world financial markets.

The U.S. economy contains elements of both cases. Because the United States is large enough to affect the world interest rate and because capital is not perfectly mobile across countries, a monetary contraction does raise the interest rate and depress investment. At the same time, a monetary contraction also raises the value of the dollar, thereby depressing net exports. Hence, although the Mundell–Fleming model does not precisely describe an economy like that of the United States, it does predict correctly what happens to international variables such as the exchange rate, and it shows how international interactions alter the effects of monetary and fiscal policies.

Summary

1. The Mundell–Fleming model is the *IS–LM* model for a small open economy. It takes the price level as given and then shows what causes fluctuations in income and the exchange rate.
2. The Mundell–Fleming model shows that fiscal policy does not influence aggregate income under floating exchange rates. A fiscal expansion causes the

currency to appreciate, reducing net exports and offsetting the usual expansionary impact on aggregate income. Fiscal policy does influence aggregate income under fixed exchange rates.

3. The Mundell–Fleming model shows that monetary policy does not influence aggregate income under fixed exchange rates. Any attempt to expand the money supply is futile, because the money supply must adjust to ensure that the exchange rate stays at its announced level. Monetary policy does influence aggregate income under floating exchange rates.
4. If investors are wary of holding assets in a country, the interest rate in that country may exceed the world interest rate by some risk premium. According to the Mundell–Fleming model, an increase in the risk premium causes the interest rate to rise and the currency of that country to depreciate.
5. There are advantages to both floating and fixed exchange rates. Floating exchange rates leave monetary policymakers free to pursue objectives other than exchange-rate stability. Fixed exchange rates reduce some of the uncertainty in international business transactions.

KEY CONCEPTS

Mundell–Fleming model	Fixed exchange rates	Revaluation
Floating exchange rates	Devaluation	

QUESTIONS FOR REVIEW

1. In the Mundell–Fleming model with floating exchange rates, explain what happens to aggregate income, the exchange rate, and the trade balance when taxes are raised. What would happen if exchange rates were fixed rather than floating?
2. In the Mundell–Fleming model with floating exchange rates, explain what happens to aggregate income, the exchange rate, and the trade balance when the money supply is reduced. What would happen if exchange rates were fixed rather than floating?
3. In the Mundell–Fleming model with floating exchange rates, explain what happens to aggregate income, the exchange rate, and the trade balance when a quota on imported cars is removed. What would happen if exchange rates were fixed rather than floating?
4. What are the advantages of floating exchange rates and fixed exchange rates?

PROBLEMS AND APPLICATIONS

1. Use the Mundell–Fleming model to predict what would happen to aggregate income, the exchange rate, and the trade balance under both floating and fixed exchange rates in response to each of the following shocks:
 - a. A fall in consumer confidence about the future induces consumers to spend less and save more.
 - b. The introduction of a stylish line of Toyotas makes some consumers prefer foreign cars over domestic cars.
 - c. The introduction of automatic teller machines reduces the demand for money.
2. The Mundell–Fleming model takes the world interest rate r^* as an exogenous variable. Let's

consider what happens when this variable changes.

- a. What might cause the world interest rate to rise?
 - b. In the Mundell–Fleming model with a floating exchange rate, what happens to aggregate income, the exchange rate, and the trade balance when the world interest rate rises?
 - c. In the Mundell–Fleming model with a fixed exchange rate, what happens to aggregate income, the exchange rate, and the trade balance when the world interest rate rises?
3. Business executives and policymakers are often concerned about the “competitiveness” of American industry (the ability of U.S. industries to sell their goods profitably in world markets).
 - a. How would a change in the exchange rate affect competitiveness?
 - b. Suppose you wanted to make domestic industries more competitive but did not want to alter aggregate income. According to the Mundell–Fleming model, what combination of monetary and fiscal policies should you pursue?
 4. Suppose that higher income implies higher imports and thus lower net exports. That is, the net-exports function is

$$NX = NX(e, Y).$$

Examine the effects in a small open economy of a fiscal expansion on income and the trade balance under

- a. A floating exchange rate.
- b. A fixed exchange rate.

How does your answer compare to the results in Table 12-1?

5. Suppose that money demand depends on disposable income, so that the equation for the money market becomes

$$M/P = L(r, Y - T).$$

Analyze the impact of a tax cut in a small open economy on the exchange rate and income under both floating and fixed exchange rates.

6. Suppose that the price level relevant for money demand includes the price of imported goods and that the price of imported goods depends on the exchange rate. That is, the money market is described by

$$M/P = L(r, Y),$$

where

$$P = \lambda P_d + (1 - \lambda)P_f/e.$$

The parameter λ is the share of domestic goods in the price index P . Assume that the price of domestic goods P_d and the price of foreign goods measured in foreign currency P_f are fixed.

- a. Suppose we graph the LM^* curve for given values of P_d and P_f (instead of the usual P). Explain why in this model this LM^* curve is upward sloping rather than vertical.
 - b. What is the effect of expansionary fiscal policy under floating exchange rates in this model? Explain. Contrast with the standard Mundell–Fleming model.
 - c. Suppose that political instability increases the country risk premium and, thereby, the interest rate. What is the effect on the exchange rate, the price level, and aggregate income in this model? Contrast with the standard Mundell–Fleming model.
7. Use the Mundell–Fleming model to answer the following questions about the state of California (a small open economy).
 - a. If California suffers from a recession, should the state government use monetary or fiscal policy to stimulate employment? Explain. (Note: For this question, assume that the state government can print dollar bills.)
 - b. If California prohibited the import of wines from the state of Washington, what would happen to income, the exchange rate, and the trade balance? Consider both the short-run and the long-run impacts.

APPENDIX

A Short-Run Model of the Large Open Economy

When analyzing policies in an economy such as the United States, we need to combine the closed-economy logic of the *IS–LM* model and the small-open-economy logic of the Mundell–Fleming model. This appendix presents a model for the intermediate case of a large open economy.

As we discussed in the appendix to Chapter 5, a large open economy differs from a small open economy because its interest rate is not fixed by world financial markets. In a large open economy, we must consider the relationship between the interest rate and the flow of capital abroad. The net capital outflow is the amount that domestic investors lend abroad minus the amount that foreign investors lend here. As the domestic interest rate falls, domestic investors find foreign lending more attractive, and foreign investors find lending here less attractive. Thus, the net capital outflow is negatively related to the interest rate. Here we add this relationship to our short-run model of national income.

The three equations of the model are

$$\begin{aligned} Y &= C(Y - T) + I(r) + G + NX(e), \\ M/P &= L(r, Y), \\ NX(e) &= CF(r). \end{aligned}$$

The first two equations are the same as those used in the Mundell–Fleming model of this chapter. The third equation, taken from the appendix to Chapter 5, states that the trade balance *NX* equals the net capital outflow *CF*, which in turn depends on the domestic interest rate.


To see what this model implies, substitute the third equation into the first, so the model becomes

$$\begin{aligned} Y &= C(Y - T) + I(r) + G + CF(r) && IS, \\ M/P &= L(r, Y) && LM. \end{aligned}$$

These two equations are much like the two equations of the closed-economy *IS–LM* model. The only difference is that expenditure now depends on the interest rate for two reasons. As before, a higher interest rate reduces investment. But now, a higher interest rate also reduces the net capital outflow and thus lowers net exports.

To analyze this model, we can use the three graphs in Figure 12-14 on page 342. Panel (a) shows the *IS–LM* diagram. As in the closed-economy model in Chapters 10 and 11, the interest rate *r* is on the vertical axis, and income *Y* is on the horizontal axis. The *IS* and *LM* curves together determine the equilibrium level of income and the equilibrium interest rate.

figure 12-14



The new net-capital-outflow term in the IS equation, $CF(r)$, makes this IS curve flatter than it would be in a closed economy. The more responsive international capital flows are to the interest rate, the flatter the IS curve is. You might recall from the Chapter 5 appendix that the small open economy represents the extreme case in which the net capital outflow is infinitely elastic at the world interest rate. In this extreme case, the IS curve is completely flat. Hence, a small open economy would be depicted in this figure with a horizontal IS curve.

Panels (b) and (c) show how the equilibrium from the $IS-LM$ model determines the net capital outflow, the trade balance, and the exchange rate. In panel (b) we see that the interest rate determines the net capital outflow. This curve slopes downward because a higher interest rate discourages domestic investors from lending abroad and encourages foreign investors to lend here. In panel (c) we see that the exchange rate adjusts to ensure that net exports of goods and services equal the net capital outflow.

Now let's use this model to examine the impact of various policies. We assume that the economy has a floating exchange rate, because this assumption is correct for most large open economies such as the United States.


Fiscal Policy

Figure 12-15 examines the impact of a fiscal expansion. An increase in government purchases or a cut in taxes shifts the IS curve to the right. As panel (a) illustrates, this shift in the IS curve leads to an increase in the level of income and an increase in the interest rate. These two effects are similar to those in a closed economy.

Yet, in the large open economy, the higher interest rate reduces the net capital outflow, as in panel (b). The fall in the net capital outflow reduces the supply of dollars in the market for foreign exchange. The exchange rate appreciates, as in panel (c). Because domestic goods become more expensive relative to foreign goods, net exports fall.

Figure 12-15 shows that a fiscal expansion does raise income in the large open economy, unlike in a small open economy under a floating exchange rate. The impact on income, however, is smaller than in a closed economy. In a closed economy, the expansionary impact of fiscal policy is partially offset by the

figure 12-15




crowding out of investment: as the interest rate rises, investment falls, reducing the fiscal-policy multipliers. In a large open economy, there is yet another offsetting factor: as the interest rate rises, the net capital outflow falls, the exchange rate appreciates, and net exports fall. Together these effects are not large enough to make fiscal policy powerless, as it is in a small open economy, but they do reduce fiscal policy's impact.

Monetary Policy

Figure 12-16 examines the effect of a monetary expansion. An increase in the money supply shifts the LM curve to the right, as in panel (a). The level of income rises, and the interest rate falls. Once again, these effects are similar to those in a closed economy.

Yet, as panel (b) shows, the lower interest rate leads to a higher net capital outflow. The increase in CF raises the supply of dollars in the market for foreign exchange. The exchange rate depreciates, as in panel (c). As domestic goods become cheaper relative to foreign goods, net exports rise.

figure 12-16



We can now see that the monetary transmission mechanism has two parts in a large open economy. As in a closed economy, a monetary expansion lowers the interest rate. As in a small open economy, a monetary expansion causes the currency to depreciate in the market for foreign exchange. The lower interest rate stimulates investment, and the lower exchange rate stimulates net exports.

A Rule of Thumb

This model of the large open economy describes well the U.S. economy today. Yet it is somewhat more complicated and cumbersome than the model of the closed economy we studied in Chapters 10 and 11 and the model of the small open economy we developed in this chapter. Fortunately, there is a useful rule of thumb to help you determine how policies influence a large open economy without remembering all the details of the model: *The large open economy is an average of the closed economy and the small open economy. To find how any policy will affect any variable, find the answer in the two extreme cases and take an average.*

For example, how does a monetary contraction affect the interest rate and investment in the short run? In a closed economy, the interest rate rises, and investment falls. In a small open economy, neither the interest rate nor investment changes. The effect in the large open economy is an average of these two cases: a monetary contraction raises the interest rate and reduces investment, but only somewhat. The fall in the net capital outflow mitigates the rise in the interest rate and the fall in investment that would occur in a closed economy. But unlike in a small open economy, the international flow of capital is not so strong as to negate fully these effects.

This rule of thumb makes the simple models all the more valuable. Although they do not describe perfectly the world in which we live, they do provide a useful guide to the effects of economic policy.

MORE PROBLEMS AND APPLICATIONS

1. Imagine that you run the central bank in a large open economy. Your goal is to stabilize income, and you adjust the money supply accordingly. Under your policy, what happens to the money supply, the interest rate, the exchange rate, and the trade balance in response to each of the following shocks?
 - a. The president raises taxes to reduce the budget deficit.
 - b. The president restricts the import of Japanese cars.
2. Over the past several decades, investors around the world have become more willing to take advantage of opportunities in other countries. Because of this increasing sophistication, economies are more open today than in the past. Consider how this development affects the ability of monetary policy to influence the economy.
 - a. If investors become more willing to substitute foreign and domestic assets, what happens to the slope of the CF function?
 - b. If the CF function changes in this way, what happens to the slope of the IS curve?
 - c. How does this change in the IS curve affect the Fed's ability to control the interest rate?
 - d. How does this change in the IS curve affect the Fed's ability to control national income?

3. Suppose that policymakers in a large open economy want to raise the level of investment without changing aggregate income or the exchange rate.
 - a. Is there any combination of domestic monetary and fiscal policies that would achieve this goal?
 - b. Is there any combination of domestic monetary, fiscal, and trade policies that would achieve this goal?
 - c. Is there any combination of monetary and fiscal policies at home and abroad that would achieve this goal?
4. Suppose that a large open economy has a fixed exchange rate.
 - a. Describe what happens in response to a fiscal contraction, such as a tax increase. Compare your answer to the case of a small open economy.
 - b. Describe what happens if the central bank expands the money supply by buying bonds from the public. Compare your answer to the case of a small open economy.

Aggregate Supply

There is always a temporary tradeoff between inflation and unemployment; there is no permanent tradeoff. The temporary tradeoff comes not from inflation per se, but from unanticipated inflation, which generally means, from a rising rate of inflation.

— Milton Friedman

Most economists analyze short-run fluctuations in aggregate income and the price level using the model of aggregate demand and aggregate supply. In the previous three chapters, we examined aggregate demand in some detail. The *IS–LM* model—together with its open-economy cousin the Mundell–Fleming model—shows how changes in monetary and fiscal policy and shocks to the money and goods markets shift the aggregate demand curve. In this chapter, we turn our attention to aggregate supply and develop theories that explain the position and slope of the aggregate supply curve.

When we introduced the aggregate supply curve in Chapter 9, we established that aggregate supply behaves differently in the short run than in the long run. In the long run, prices are flexible, and the aggregate supply curve is vertical. When the aggregate supply curve is vertical, shifts in the aggregate demand curve affect the price level, but the output of the economy remains at its natural rate. By contrast, in the short run, prices are sticky, and the aggregate supply curve is not vertical. In this case, shifts in aggregate demand do cause fluctuations in output. In Chapter 9 we took a simplified view of price stickiness by drawing the short-run aggregate supply curve as a horizontal line, representing the extreme situation in which all prices are fixed. Our task now is to refine this understanding of short-run aggregate supply.

Unfortunately, one fact makes this task more difficult: economists disagree about how best to explain aggregate supply. As a result, this chapter begins by presenting three prominent models of the short-run aggregate supply curve. Among economists, each of these models has some prominent adherents (as well as some prominent critics), and you can decide for yourself which you find most plausible. Although these models differ in some significant details, they are also related in an important way: they share a common theme about what makes the

short-run and long-run aggregate supply curves differ and a common conclusion that the short-run aggregate supply curve is upward sloping.

After examining the models, we examine an implication of the short-run aggregate supply curve. We show that this curve implies a tradeoff between two measures of economic performance—inflation and unemployment. According to this tradeoff, to reduce the rate of inflation policymakers must temporarily raise unemployment, and to reduce unemployment they must accept higher inflation. As the quotation at the beginning of the chapter suggests, the tradeoff between inflation and unemployment is only temporary. One goal of this chapter is to explain why policymakers face such a tradeoff in the short run and, just as important, why they do not face it in the long run.

13-1 Three Models of Aggregate Supply

When classes in physics study balls rolling down inclined planes, they often begin by assuming away the existence of friction. This assumption makes the problem simpler and is useful in many circumstances, but no good engineer would ever take this assumption as a literal description of how the world works. Similarly, this book began with classical macroeconomic theory, but it would be a mistake to assume that this model is always true. Our job now is to look more deeply into the “frictions” of macroeconomics.

We do this by examining three prominent models of aggregate supply, roughly in the order of their development. In all the models, some market imperfection (that is, some type of friction) causes the output of the economy to deviate from the classical benchmark. As a result, the short-run aggregate supply curve is upward sloping, rather than vertical, and shifts in the aggregate demand curve cause the level of output to deviate temporarily from the natural rate. These temporary deviations represent the booms and busts of the business cycle.

Although each of the three models takes us down a different theoretical route, each route ends up in the same place. That final destination is a short-run aggregate supply equation of the form

$$Y = \bar{Y} + \alpha(P - P^e), \quad \alpha > 0$$

where Y is output, $\bar{Y}$ is the natural rate of output, P is the price level, and P^e is the expected price level. This equation states that output deviates from its natural rate when the price level deviates from the expected price level. The parameter α indicates how much output responds to unexpected changes in the price level; $1/\alpha$ is the slope of the aggregate supply curve.

Each of the three models tells a different story about what lies behind this short-run aggregate supply equation. In other words, each highlights a particular reason why unexpected movements in the price level are associated with fluctuations in aggregate output.

The Sticky-Wage Model

To explain why the short-run aggregate supply curve is upward sloping, many economists stress the sluggish adjustment of nominal wages. In many industries, nominal wages are set by long-term contracts, so wages cannot adjust quickly when economic conditions change. Even in industries not covered by formal contracts, implicit agreements between workers and firms may limit wage changes. Wages may also depend on social norms and notions of fairness that evolve slowly. For these reasons, many economists believe that nominal wages are sticky in the short run.

The **sticky-wage model** shows what a sticky nominal wage implies for aggregate supply. To preview the model, consider what happens to the amount of output produced when the price level rises:

1. When the nominal wage is stuck, a rise in the price level lowers the real wage, making labor cheaper.
2. The lower real wage induces firms to hire more labor.
3. The additional labor hired produces more output.

This positive relationship between the price level and the amount of output means that the aggregate supply curve slopes upward during the time when the nominal wage cannot adjust.

To develop this story of aggregate supply more formally, assume that workers and firms bargain over and agree on the nominal wage before they know what the price level will be when their agreement takes effect. The bargaining parties—the workers and the firms—have in mind a target real wage. The target may be the real wage that equilibrates labor supply and demand. More likely, the target real wage is higher than the equilibrium real wage: as discussed in Chapter 6, union power and efficiency-wage considerations tend to keep real wages above the level that brings supply and demand into balance.

The workers and firms set the nominal wage W based on the target real wage ω and on their expectation of the price level P^e . The nominal wage they set is

$$W = \omega \times P^e$$

Nominal Wage = Target Real Wage $\times$ Expected Price Level.

After the nominal wage has been set and before labor has been hired, firms learn the actual price level P . The real wage turns out to be

$$W/P = \omega \times (P^e/P)$$

Real Wage = Target Real Wage $\times \frac{\text{Expected Price Level}}{\text{Actual Price Level}}$.

This equation shows that the real wage deviates from its target if the actual price level differs from the expected price level. When the actual price level is greater than expected, the real wage is less than its target; when the actual price level is less than expected, the real wage is greater than its target.

The final assumption of the sticky-wage model is that employment is determined by the quantity of labor that firms demand. In other words, the bargain between the workers and the firms does not determine the level of employment in advance; instead, the workers agree to provide as much labor as the firms wish to buy at the predetermined wage. We describe the firms' hiring decisions by the labor demand function

$$L = L^d(W/P),$$


which states that the lower the real wage, the more labor firms hire. The labor demand curve is shown in panel (a) of Figure 13-1. Output is determined by the production function

$$Y = F(L),$$

which states that the more labor is hired, the more output is produced. This is shown in panel (b) of Figure 13-1.

Panel (c) of Figure 13-1 shows the resulting aggregate supply curve. Because the nominal wage is sticky, an unexpected change in the price level

figure 13-1



moves the real wage away from the target real wage, and this change in the real wage influences the amounts of labor hired and output produced. The aggregate supply curve can be written as

$$Y = \bar{Y} + \alpha(P - P^e).$$

Output deviates from its natural level when the price level deviates from the expected price level.¹

CASE STUDY

The Cyclical Behavior of the Real Wage

In any model with an unchanging labor demand curve, such as the model we just discussed, employment rises when the real wage falls. In the sticky-wage model, an unexpected rise in the price level lowers the real wage and thereby raises the quantity of labor hired and the amount of output produced. Thus, the real wage should be *countercyclical*: it should fluctuate in the opposite direction from employment and output. Keynes himself wrote in *The General Theory* that “an increase in employment can only occur to the accompaniment of a decline in the rate of real wages.”


The earliest attacks on *The General Theory* came from economists challenging Keynes’s prediction. Figure 13–2 is a scatterplot of the percentage change in real compensation per hour and the percentage change in real GDP using annual data for the U.S. economy from 1960 to 2000. If Keynes’s prediction were correct, the dots in this figure would show a downward-sloping pattern, indicating a negative relationship. Yet the figure shows only a weak correlation between the real wage and output, and it is the opposite of what Keynes predicted. That is, if the real wage is cyclical at all, it is slightly *procyclical*: the real wage tends to rise when output rises. Abnormally high labor costs cannot explain the low employment and output observed in recessions.

How should we interpret this evidence? Most economists conclude that the sticky-wage model cannot fully explain aggregate supply. They advocate models in which the labor demand curve shifts over the business cycle. These shifts may arise because firms have sticky prices and cannot sell all they want at those prices; we discuss this possibility later. Alternatively, the labor demand curve may shift because of shocks to technology, which alter labor productivity. The theory we discuss in Chapter 19, called the theory of real business cycles, gives a prominent role to technology shocks as a source of economic fluctuations.²

¹ For more on the sticky-wage model, see Jo Anna Gray, “Wage Indexation: A Macroeconomic Approach,” *Journal of Monetary Economics* 2 (April 1976): 221–235; and Stanley Fischer, “Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule,” *Journal of Political Economy* 85 (February 1977): 191–205.

² For some of the recent work on the cyclical behavior of the real wage, see Scott Sumner and Stephen Silver, “Real Wages, Employment, and the Phillips Curve,” *Journal of Political Economy* 97 (June 1989): 706–720; and Gary Solon, Robert Barsky, and Jonathan A. Parker, “Measuring the Cyclical Behavior of Real Wages: How Important Is Composition Bias?” *Quarterly Journal of Economics* 109 (February 1994): 1–25.

figure 13-2



The Cyclical Behavior of the Real Wage This scatterplot shows the percentage change in real GDP and the percentage change in the real wage (measured here as real private hourly earnings). As output fluctuates, the real wage typically moves in the same direction. That is, the real wage is somewhat procyclical. This observation is inconsistent with the sticky-wage model.

Source: U.S. Department of Commerce and U.S. Department of Labor.

The Imperfect-Information Model

The second explanation for the upward slope of the short-run aggregate supply curve is called the **imperfect-information model**. Unlike the sticky-wage model, this model assumes that markets clear—that is, all wages and prices are free to adjust to balance supply and demand. In this model, the short-run and long-run aggregate supply curves differ because of temporary misperceptions about prices.

The imperfect-information model assumes that each supplier in the economy produces a single good and consumes many goods. Because the number of goods is so large, suppliers cannot observe all prices at all times. They monitor closely the prices of what they produce but less closely the prices of all the goods they consume. Because of imperfect information, they sometimes confuse changes in the overall level of prices with changes in relative prices. This confusion influences decisions about how much to supply, and it leads to a positive relationship between the price level and output in the short run.

Consider the decision facing a single supplier—a wheat farmer, for instance. Because the farmer earns income from selling wheat and uses this income to buy goods and services, the amount of wheat she chooses to produce depends on the

price of wheat relative to the prices of other goods and services in the economy. If the relative price of wheat is high, the farmer is motivated to work hard and produce more wheat, because the reward is great. If the relative price of wheat is low, she prefers to enjoy more leisure and produce less wheat.

Unfortunately, when the farmer makes her production decision, she does not know the relative price of wheat. As a wheat producer, she monitors the wheat market closely and always knows the nominal price of wheat. But she does not know the prices of all the other goods in the economy. She must, therefore, estimate the relative price of wheat using the nominal price of wheat and her expectation of the overall price level.

Consider how the farmer responds if all prices in the economy, including the price of wheat, increase. One possibility is that she expected this change in prices. When she observes an increase in the price of wheat, her estimate of its relative price is unchanged. She does not work any harder.

The other possibility is that the farmer did not expect the price level to increase (or to increase by this much). When she observes the increase in the price of wheat, she is not sure whether other prices have risen (in which case wheat's relative price is unchanged) or whether only the price of wheat has risen (in which case its relative price is higher). The rational inference is that some of each has happened. In other words, the farmer infers from the increase in the nominal price of wheat that its relative price has risen somewhat. She works harder and produces more.

Our wheat farmer is not unique. When the price level rises unexpectedly, all suppliers in the economy observe increases in the prices of the goods they produce. They all infer, rationally but mistakenly, that the relative prices of the goods they produce have risen. They work harder and produce more.

To sum up, the imperfect-information model says that when actual prices exceed expected prices, suppliers raise their output. The model implies an aggregate supply curve that is now familiar:

$$Y = \bar{Y} + \alpha(P - P^e).$$

Output deviates from the natural rate when the price level deviates from the expected price level.³

The Sticky-Price Model

Our third explanation for the upward-sloping short-run aggregate supply curve is called the **sticky-price model**. This model emphasizes that firms do not instantly adjust the prices they charge in response to changes in demand. Sometimes prices are set by long-term contracts between firms and customers. Even

³ Two economists who have emphasized the role of imperfect information for understanding the short-run effects of monetary policy are the Nobel Prize winners Milton Friedman and Robert Lucas. See Milton Friedman, "The Role of Monetary Policy," *American Economic Review* 58 (March 1968): 1–17; and Robert E. Lucas, Jr., "Understanding Business Cycles," *Stabilization of the Domestic and International Economy*, vol. 5 of Carnegie-Rochester Conference on Public Policy (Amsterdam: North-Holland, 1977).

without formal agreements, firms may hold prices steady in order not to annoy their regular customers with frequent price changes. Some prices are sticky because of the way markets are structured: once a firm has printed and distributed its catalog or price list, it is costly to alter prices.

To see how sticky prices can help explain an upward-sloping aggregate supply curve, we first consider the pricing decisions of individual firms and then add together the decisions of many firms to explain the behavior of the economy as a whole. Notice that this model encourages us to depart from the assumption of perfect competition, which we have used since Chapter 3. Perfectly competitive firms are price takers rather than price setters. If we want to consider how firms set prices, it is natural to assume that these firms have at least some monopoly control over the prices they charge.

Consider the pricing decision facing a typical firm. The firm's desired price p depends on two macroeconomic variables:

- The overall level of prices P . A higher price level implies that the firm's costs are higher. Hence, the higher the overall price level, the more the firm would like to charge for its product.
- The level of aggregate income Y . A higher level of income raises the demand for the firm's product. Because marginal cost increases at higher levels of production, the greater the demand, the higher the firm's desired price.

We write the firm's desired price as

$$p = P + a(Y - \bar{Y}).$$

This equation says that the desired price p depends on the overall level of prices P and on the level of aggregate output relative to the natural rate $Y - \bar{Y}$. The parameter a (which is greater than zero) measures how much the firm's desired price responds to the level of aggregate output.⁴

Now assume that there are two types of firms. Some have flexible prices: they always set their prices according to this equation. Others have sticky prices: they announce their prices in advance based on what they expect economic conditions to be. Firms with sticky prices set prices according to

$$p = P^e + a(Y^e - \bar{Y}^e),$$

where, as before, a superscript "e" represents the expected value of a variable. For simplicity, assume that these firms expect output to be at its natural rate, so that the last term, $a(Y^e - \bar{Y}^e)$, is zero. Then these firms set the price

$$p = P^e.$$

That is, firms with sticky prices set their prices based on what they expect other firms to charge.

⁴ *Mathematical note:* The firm cares most about its relative price, which is the ratio of its nominal price to the overall price level. If we interpret p and P as the logarithms of the firm's price and the price level, then this equation states that the desired relative price depends on the deviation of output from the natural rate.

We can use the pricing rules of the two groups of firms to derive the aggregate supply equation. To do this, we find the overall price level in the economy, which is the weighted average of the prices set by the two groups. If s is the fraction of firms with sticky prices and $1 - s$ the fraction with flexible prices, then the overall price level is

$$P = sP^e + (1 - s)[P + a(Y - \bar{Y})].$$

The first term is the price of the sticky-price firms weighted by their fraction in the economy, and the second term is the price of the flexible-price firms weighted by their fraction. Now subtract $(1 - s)P$ from both sides of this equation to obtain

$$sP = sP^e + (1 - s)[a(Y - \bar{Y})].$$

Divide both sides by s to solve for the overall price level:

$$P = P^e + [(1 - s)a/s](Y - \bar{Y}).$$

The two terms in this equation are explained as follows:

- When firms expect a high price level, they expect high costs. Those firms that fix prices in advance set their prices high. These high prices cause the other firms to set high prices also. Hence, a high expected price level P^e leads to a high actual price level P .
- When output is high, the demand for goods is high. Those firms with flexible prices set their prices high, which leads to a high price level. The effect of output on the price level depends on the proportion of firms with flexible prices.

Hence, the overall price level depends on the expected price level and on the level of output.

Algebraic rearrangement puts this aggregate pricing equation into a more familiar form:

$$Y = \bar{Y} + \alpha(P - P^e),$$

where $\alpha = s/[(1 - s)a]$. Like the other models, the sticky-price model says that the deviation of output from the natural rate is positively associated with the deviation of the price level from the expected price level.

Although the sticky-price model emphasizes the goods market, consider briefly what is happening in the labor market. If a firm's price is stuck in the short run, then a reduction in aggregate demand reduces the amount that the firm is able to sell. The firm responds to the drop in sales by reducing its production and its demand for labor. Note the contrast to the sticky-wage model: the firm here does not move along a fixed labor demand curve. Instead, fluctuations in output are associated with shifts in the labor demand curve. Because of these shifts in labor demand, employment, production, and the real wage can all move in the same direction. Thus, the real wage can be procyclical.⁵

⁵ For a more advanced development of the sticky-price model, see Julio Rotemberg, "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies* 49 (1982): 517–531.

CASE STUDY

International Differences in the Aggregate Supply Curve

Although all countries experience economic fluctuations, these fluctuations are not exactly the same everywhere. International differences are intriguing puzzles in themselves, and they often provide a way to test alternative economic theories. Examining international differences has been especially fruitful in research on aggregate supply.

When economist Robert Lucas proposed the imperfect-information model, he derived a surprising interaction between aggregate demand and aggregate supply: according to his model, the slope of the aggregate supply curve should depend on the volatility of aggregate demand. In countries where aggregate demand fluctuates widely, the aggregate price level fluctuates widely as well. Because most movements in prices in these countries do not represent movements in relative prices, suppliers should have learned not to respond much to unexpected changes in the price level. Therefore, the aggregate supply curve should be relatively steep (that is, α will be small). Conversely, in countries where aggregate demand is relatively stable, suppliers should have learned that most price changes are relative price changes. Accordingly, in these countries, suppliers should be more responsive to unexpected price changes, making the aggregate supply curve relatively flat (that is, α will be large).

Lucas tested this prediction by examining international data on output and prices. He found that changes in aggregate demand have the biggest effect on output in those countries where aggregate demand and prices are most stable. Lucas concluded that the evidence supports the imperfect-information model.⁶

The sticky-price model also makes predictions about the slope of the short-run aggregate supply curve. In particular, it predicts that the average rate of inflation should influence the slope of the short-run aggregate supply curve. When the average rate of inflation is high, it is very costly for firms to keep prices fixed for long intervals. Thus, firms adjust prices more frequently. More frequent price adjustment in turn allows the overall price level to respond more quickly to shocks to aggregate demand. Hence, a high rate of inflation should make the short-run aggregate supply curve steeper.

International data support this prediction of the sticky-price model. In countries with low average inflation, the short-run aggregate supply curve is relatively flat: fluctuations in aggregate demand have large effects on output and are slowly reflected in prices. High-inflation countries have steep short-run aggregate supply curves. In other words, high inflation appears to erode the frictions that cause prices to be sticky.⁷

Note that the sticky-price model can also explain Lucas's finding that countries with variable aggregate demand have steep aggregate supply curves. If the price level is highly variable, few firms will commit to prices in advance (s will be small). Hence, the aggregate supply curve will be steep (α will be small).

⁶ Robert E. Lucas, Jr., "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review* 63 (June 1973): 326–334.

⁷ Laurence Ball, N. Gregory Mankiw, and David Romer, "The New Keynesian Economics and the Output-Inflation Tradeoff," *Brookings Papers on Economic Activity* (1988:1): 1–65.

Summary and Implications

We have seen three models of aggregate supply and the market imperfection that each uses to explain why the short-run aggregate supply curve is upward sloping. One model assumes nominal wages are sticky; the second assumes information about prices is imperfect; the third assumes prices are sticky. Keep in mind that these models are not incompatible with one another. We need not accept one model and reject the others. The world may contain all three of these market imperfections, and all may contribute to the behavior of short-run aggregate supply.


Although the three models of aggregate supply differ in their assumptions and emphases, their implications for aggregate output are similar. All can be summarized by the equation

$$Y = \bar{Y} + \alpha(P - P^e).$$

This equation states that deviations of output from the natural rate are related to deviations of the price level from the expected price level. *If the price level is higher than the expected price level, output exceeds its natural rate. If the price level is lower than the expected price level, output falls short of its natural rate.* Figure 13-3 graphs this equation. Notice that the short-run aggregate supply curve is drawn for a given expectation P^e and that a change in P^e would shift the curve.


Now that we have a better understanding of aggregate supply, let's put aggregate supply and aggregate demand back together. Figure 13-4 uses our aggregate supply equation to show how the economy responds to an unexpected increase in aggregate demand attributable, say, to an unexpected monetary expansion. In the short run, the equilibrium moves from point A to point B. The increase in aggregate demand raises the actual price level from P_1 to P_2 . Because people did not expect this increase in the price level, the expected price level remains at P_1^e , and output rises from Y_1 to Y_2 , which is above the natural rate $\bar{Y}$. Thus, the unexpected expansion in aggregate demand causes the economy to boom.

figure 13-3



The Short-Run Aggregate Supply Curve Output deviates from the natural rate $\bar{Y}$ if the price level P deviates from the expected price level P^e .

figure 13-4



How Shifts in Aggregate Demand Lead to Short-Run Fluctuations Here the economy begins in a long-run equilibrium, point A. When aggregate demand increases unexpectedly, the price level rises from P_1 to P_2 . Because the price level P_2 is above the expected price level P_2^e , output rises temporarily above the natural rate, as the economy moves along the short-run aggregate supply curve from point A to point B. In the long run, the expected price level rises to P_3^e , causing the short-run aggregate supply curve to shift upward. The economy returns to a new long-run equilibrium, point C, where output is back at its natural rate.

Yet the boom does not last forever. In the long run, the expected price level rises to catch up with reality, causing the short-run aggregate supply curve to shift upward. As the expected price level rises from P_2^e to P_3^e , the equilibrium of the economy moves from point B to point C. The actual price level rises from P_2 to P_3 , and output falls from Y_2 to Y_3 . In other words, the economy returns to the natural level of output in the long run, but at a much higher price level.

This analysis shows an important principle, which holds for each of the three models of aggregate supply: long-run monetary neutrality and short-run monetary nonneutrality are perfectly compatible. Short-run nonneutrality is represented here by the movement from point A to point B, and long-run monetary neutrality is represented by the movement from point A to point C. We reconcile the short-run and long-run effects of money by emphasizing the adjustment of expectations about the price level.

13-2 Inflation, Unemployment, and the Phillips Curve

Two goals of economic policymakers are low inflation and low unemployment, but often these goals conflict. Suppose, for instance, that policymakers were to use monetary or fiscal policy to expand aggregate demand. This policy would move the economy along the short-run aggregate supply curve to a point of higher output and a higher price level. (Figure 13-4 shows this as the change from point A to point B.) Higher output means lower unemployment, because firms need more workers when they produce more. A higher price level, given

the previous year's price level, means higher inflation. Thus, when policymakers move the economy up along the short-run aggregate supply curve, they reduce the unemployment rate and raise the inflation rate. Conversely, when they contract aggregate demand and move the economy down the short-run aggregate supply curve, unemployment rises and inflation falls.

This tradeoff between inflation and unemployment, called the *Phillips curve*, is our topic in this section. As we have just seen (and will derive more formally in a moment), the Phillips curve is a reflection of the short-run aggregate supply curve: as policymakers move the economy along the short-run aggregate supply curve, unemployment and inflation move in opposite directions. The Phillips curve is a useful way to express aggregate supply because inflation and unemployment are such important measures of economic performance.

Deriving the Phillips Curve From the Aggregate Supply Curve

The **Phillips curve** in its modern form states that the inflation rate depends on three forces:

- Expected inflation;
- The deviation of unemployment from the natural rate, called *cyclical unemployment*;
- Supply shocks.

These three forces are expressed in the following equation:

$$\pi = \pi^e - \beta(u - u^n) + v$$

$$\text{Inflation} = \text{Expected Inflation} - (\beta \times \text{Cyclical Unemployment}) + \text{Supply Shock},$$

where β is a parameter measuring the response of inflation to cyclical unemployment. Notice that there is a minus sign before the cyclical unemployment term: high unemployment tends to reduce inflation. This equation summarizes the relationship between inflation and unemployment.

From where does this equation for the Phillips curve come? Although it may not seem familiar, we can derive it from our equation for aggregate supply. To see how, write the aggregate supply equation as

$$P = P^e + (1/\alpha)(Y - \bar{Y}).$$

With one addition, one subtraction, and one substitution, we can manipulate this equation to yield a relationship between inflation and unemployment.

Here are the three steps. First, add to the right-hand side of the equation a supply shock v to represent exogenous events (such as a change in world oil prices) that alter the price level and shift the short-run aggregate supply curve:

$$P = P^e + (1/\alpha)(Y - \bar{Y}) + v.$$

Next, to go from the price level to inflation rates, subtract last year's price level P_{-1} from both sides of the equation to obtain

$$(P - P_{-1}) = (P^e - P_{-1}) + (1/\alpha)(Y - \bar{Y}) + v.$$

The term on the left-hand side, $P - P_{-1}$, is the difference between the current price level and last year's price level, which is inflation π .⁸ The term on the right-hand side, $P^e - P_{-1}$, is the difference between the expected price level and last year's price level, which is expected inflation π^e . Therefore, we can replace $P - P_{-1}$ with π and $P^e - P_{-1}$ with π^e :

$$\pi = \pi^e + (1/\alpha)(Y - \bar{Y}) + v.$$

Third, to go from output to unemployment, recall from Chapter 2 that Okun's law gives a relationship between these two variables. One version of Okun's law states that the deviation of output from its natural rate is inversely related to the deviation of unemployment from its natural rate; that is, when output is higher than the natural rate of output, unemployment is lower than the natural rate of unemployment. We can write this as

$$(1/\alpha)(Y - \bar{Y}) = -\beta(u - u^n).$$

Using this Okun's law relationship, we can substitute $-\beta(u - u^n)$ for $(1/\alpha)(Y - \bar{Y})$ in the previous equation to obtain

$$\pi = \pi^e - \beta(u - u^n) + v.$$

Thus, we can derive the Phillips curve equation from the aggregate supply equation.

All this algebra is meant to show one thing: the Phillips curve equation and the short-run aggregate supply equation represent essentially the same macroeconomic ideas. In particular, both equations show a link between real and nominal variables that causes the classical dichotomy (the theoretical separation of real and nominal variables) to break down in the short run. According to the short-run aggregate supply equation, output is related to unexpected movements in the price level. According to the Phillips curve equation, unemployment is related to unexpected movements in the inflation rate. The aggregate supply curve is more convenient when we are studying output and the price level, whereas the Phillips curve is more convenient when we are studying unemployment and inflation. But we should not lose sight of the fact that the Phillips curve and the aggregate supply curve are two sides of the same coin.

⁸ *Mathematical note:* This statement is not precise, because inflation is really the *percentage* change in the price level. To make the statement more precise, interpret P as the logarithm of the price level. By the properties of logarithms, the change in P is roughly the inflation rate. The reason is that $dP = d(\log \text{ price level}) = d(\text{price level})/\text{price level}$.



The History of the Modern Phillips Curve

The Phillips curve is named after New Zealand-born economist A. W. Phillips. In 1958 Phillips observed a negative relationship between the unemployment rate and the rate of wage inflation in data for the United Kingdom.⁹ The Phillips curve that economists use today differs in three ways from the relationship Phillips examined.

First, the modern Phillips curve substitutes price inflation for wage inflation. This difference is not crucial, because price inflation and wage inflation are closely related. In periods when wages are rising quickly, prices are rising quickly as well.

Second, the modern Phillips curve includes expected inflation. This addition is due to the work of Milton Friedman and Edmund Phelps. In developing early versions of the imperfect information model in the 1960s, these two economists emphasized the importance of expectations for aggregate supply.

Third, the modern Phillips curve includes supply shocks. Credit for this addition goes to OPEC, the Organization of Petroleum Exporting Countries. In the 1970s OPEC caused large increases in the world price of oil, which made economists more aware of the importance of shocks to aggregate supply.

Adaptive Expectations and Inflation Inertia

To make the Phillips curve useful for analyzing the choices facing policymakers, we need to say what determines expected inflation. A simple and often plausible assumption is that people form their expectations of inflation based on recently observed inflation. This assumption is called **adaptive expectations**. For example, suppose that people expect prices to rise this year at the same rate as they did last year. Then expected inflation π^e equals last year's inflation π_{-1} :

$$\pi^e = \pi_{-1}.$$

In this case, we can write the Phillips curve as

$$\pi = \pi_{-1} - \beta(u - u^n) + v,$$

which states that inflation depends on past inflation, cyclical unemployment, and a supply shock. When the Phillips curve is written in this form, the natural rate of unemployment is sometimes called the Non-Accelerating Inflation Rate of Unemployment, or *NAIRU*.

The first term in this form of the Phillips curve, π_{-1} , implies that inflation has inertia. That is, like an object moving through space, inflation keeps going unless something acts to stop it. In particular, if unemployment is at the NAIRU and if there are no supply shocks, the continued rise in price level neither speeds up nor slows down. This inertia arises because past inflation influences expectations of future inflation and because these expectations influence the wages and prices

⁹ A. W. Phillips, "The Relationship Between Unemployment and the Rate of Change of Money Wages in the United Kingdom, 1861–1957," *Economica* 25 (November 1958): 283–299.

that people set. Robert Solow captured the concept of inflation inertia well when, during the high inflation of the 1970s, he wrote, “Why is our money ever less valuable? Perhaps it is simply that we have inflation because we expect inflation, and we expect inflation because we’ve had it.”

In the model of aggregate supply and aggregate demand, inflation inertia is interpreted as persistent upward shifts in both the aggregate supply curve and the aggregate demand curve. Consider first aggregate supply. If prices have been rising quickly, people will expect them to continue to rise quickly. Because the position of the short-run aggregate supply curve depends on the expected price level, the short-run aggregate supply curve will shift upward over time. It will continue to shift upward until some event, such as a recession or a supply shock, changes inflation and thereby changes expectations of inflation.

The aggregate demand curve must also shift upward to confirm the expectations of inflation. Most often, the continued rise in aggregate demand is caused by persistent growth in the money supply. If the Fed suddenly halted money growth, aggregate demand would stabilize, and the upward shift in aggregate supply would cause a recession. The high unemployment in the recession would reduce inflation and expected inflation, causing inflation inertia to subside.

Two Causes of Rising and Falling Inflation

The second and third terms in the Phillips curve equation show the two forces that can change the rate of inflation.

The second term, $\beta(u - u^n)$, shows that cyclical unemployment—the deviation of unemployment from its natural rate—exerts upward or downward pressure on inflation. Low unemployment pulls the inflation rate up. This is called **demand-pull inflation** because high aggregate demand is responsible for this type of inflation. High unemployment pulls the inflation rate down. The parameter β measures how responsive inflation is to cyclical unemployment.


The third term, v , shows that inflation also rises and falls because of supply shocks. An adverse supply shock, such as the rise in world oil prices in the 1970s, implies a positive value of v and causes inflation to rise. This is called **cost-push inflation** because adverse supply shocks are typically events that push up the costs of production. A beneficial supply shock, such as the oil glut that led to a fall in oil prices in the 1980s, makes v negative and causes inflation to fall.

CASE STUDY

Inflation and Unemployment in the United States

Because inflation and unemployment are such important measures of economic performance, macroeconomic developments are often viewed through the lens of the Phillips curve. Figure 13-5 displays the history of inflation and unemployment in the United States since 1961. These four decades of data illustrate some of the causes of rising or falling inflation.

figure 13-5



Inflation and Unemployment in the United States Since 1961 This figure uses annual data on the unemployment rate and the inflation rate (percentage change in the GDP deflator) to illustrate macroeconomic developments over the past four decades.

Source: U.S. Department of Commerce and U.S. Department of Labor.

The 1960s showed how policymakers can, in the short run, lower unemployment at the cost of higher inflation. The tax cut of 1964, together with expansionary monetary policy, expanded aggregate demand and pushed the unemployment rate below 5 percent. This expansion of aggregate demand continued in the late 1960s largely as a by-product of government spending for the Vietnam War. Unemployment fell lower and inflation rose higher than policymakers intended.

The 1970s were a period of economic turmoil. The decade began with policymakers trying to lower the inflation inherited from the 1960s. President Nixon imposed temporary controls on wages and prices, and the Federal Reserve engineered a recession through contractionary monetary policy, but the inflation rate fell only slightly. The effects of wage and price controls ended when the controls were lifted, and the recession was too small to counteract the inflationary impact of the boom that had preceded it. By 1972 the unemployment rate was the same as a decade earlier, whereas inflation was 3 percentage points higher.

Beginning in 1973 policymakers had to cope with the large supply shocks caused by the Organization of Petroleum Exporting Countries (OPEC). OPEC first raised oil prices in the mid-1970s, pushing the inflation rate up to about 10 percent. This adverse supply shock, together with temporarily tight monetary policy, led to a recession in 1975. High unemployment during the recession reduced inflation somewhat, but further OPEC price hikes pushed inflation up again in the late 1970s.

The 1980s began with high inflation and high expectations of inflation. Under the leadership of Chairman Paul Volcker, the Federal Reserve doggedly pursued monetary policies aimed at reducing inflation. In 1982 and 1983 the unemployment rate reached its highest level in 40 years. High unemployment, aided by a fall in oil prices in 1986, pulled the inflation rate down from about 10 percent to about 3 percent. By 1987 the unemployment rate of about 6 percent was close to most estimates of the natural rate. Unemployment continued to fall through the 1980s, however, reaching a low of 5.2 percent in 1989 and beginning a new round of demand-pull inflation.

Compared to the previous 30 years, the 1990s were relatively quiet. The decade began with a recession caused by several contractionary shocks to aggregate demand: tight monetary policy, the savings-and-loan crisis, and a fall in consumer confidence coinciding with the Gulf War. The unemployment rate rose to 7.3 percent in 1992. Inflation fell, but only slightly. Unlike in the 1982 recession, unemployment in the 1990 recession was never far above the natural rate, so the effect on inflation was small.

By the late 1990s, inflation and unemployment both reached their lowest levels in many years. Some economists explain this fortunate development by claiming that the economy's natural rate of unemployment fell (for reasons discussed in Chapter 6). Others argue that various temporary factors (such as a strong U.S. dollar attributable to a financial crisis in Asia) yielded favorable supply shocks. Most likely, a combination of events helped keep inflation in check, despite low unemployment. In 2000, however, inflation did begin to creep up.


Thus, U.S. macroeconomic history exhibits the many causes of inflation. The 1960s and the 1980s show the two sides of demand-pull inflation: in the 1960s low unemployment pulled inflation up, and in the 1980s high unemployment pulled inflation down. The 1970s with their oil-price hikes show the effects of cost-push inflation.

The Short-Run Tradeoff Between Inflation and Unemployment

Consider the options the Phillips curve gives to a policymaker who can influence aggregate demand with monetary or fiscal policy. At any moment, expected inflation and supply shocks are beyond the policymaker's immediate control. Yet, by changing aggregate demand, the policymaker can alter output, unemployment, and inflation. The policymaker can expand aggregate demand to lower unemployment and raise inflation. Or the policymaker can depress aggregate demand to raise unemployment and lower inflation.

Figure 13-6 plots the Phillips curve equation and shows the short-run tradeoff between inflation and unemployment. When unemployment is at its natural rate ($u = u^n$), inflation depends on expected inflation and the supply shock ($\pi = \pi^e + v$). The parameter β determines the slope of the tradeoff between inflation and unemployment. In the short run, for a given level of expected inflation, policymakers can manipulate aggregate demand to choose a

figure 13-6




The Short-Run Tradeoff Between Inflation and Unemployment In the short run, inflation and unemployment are negatively related. At any point in time, a policymaker who controls aggregate demand can choose a combination of inflation and unemployment on this short-run Phillips curve.

combination of inflation and unemployment on this curve, called the *short-run Phillips curve*.

Notice that the position of the short-run Phillips curve depends on the expected rate of inflation. If expected inflation rises, the curve shifts upward, and the policymaker's tradeoff becomes less favorable: inflation is higher for any level of unemployment. Figure 13-7 shows how the tradeoff depends on expected inflation.

Because people adjust their expectations of inflation over time, the tradeoff between inflation and unemployment holds only in the short run. The policymaker cannot keep inflation above expected inflation (and thus unemployment below its natural rate) forever. Eventually, expectations adapt to whatever inflation rate the

figure 13-7



Shifts in the Short-Run Tradeoff The short-run tradeoff between inflation and unemployment depends on expected inflation. The curve is higher when expected inflation is higher.

FYI

How Precise Are Estimates of the Natural Rate of Unemployment?

If you ask an astronomer how far a particular star is from our sun, he'll give you a number, but it won't be accurate. Man's ability to measure astronomical distances is still limited. An astronomer might well take better measurements and conclude that a star is really twice or half as far away as he previously thought.

Estimates of the natural rate of unemployment, or NAIRU, are also far from precise. One problem is supply shocks. Shocks to oil supplies, farm harvests, or technological progress can cause inflation to rise or fall in the short run. When we observe rising inflation, therefore, we cannot be sure if it is evidence that the unemployment rate is below the natural rate or evidence that the economy is experiencing an adverse supply shock.

A second problem is that the natural rate changes over time. Demographic changes (such as the aging of the baby-boom generation), policy changes (such as minimum-wage laws), and institutional changes (such as the declining role

of unions) all influence the economy's normal level of unemployment. Estimating the natural rate is like hitting a moving target.

Economists deal with these problems using statistical techniques that yield a best guess about the natural rate and allow them to gauge the uncertainty associated with their estimates. In one such study, Douglas Staiger, James Stock, and Mark Watson estimated the natural rate to be 6.2 percent in 1990, with a 95-percent confidence interval from 5.1 to 7.7 percent. A 95-percent confidence interval is a range such that the statistician is 95-percent confident that the true value falls in that range. The large confidence interval here of 2.6 percentage points shows that estimates of the natural rate are not at all precise.

This conclusion has profound implications. Policymakers may want to keep unemployment close to its natural rate, but their ability to do so is limited by the fact that we cannot be sure what that natural rate is.¹⁰

policymaker has chosen. In the long run, the classical dichotomy holds, unemployment returns to its natural rate, and there is no tradeoff between inflation and unemployment.

Disinflation and the Sacrifice Ratio

Imagine an economy in which unemployment is at its natural rate and inflation is running at 6 percent. What would happen to unemployment and output if the central bank pursued a policy to reduce inflation from 6 to 2 percent?

The Phillips curve shows that in the absence of a beneficial supply shock, lowering inflation requires a period of high unemployment and reduced output. But by how much and for how long would unemployment need to rise above the natural rate? Before deciding whether to reduce inflation, policymakers must know how much output would be lost during the transition to lower inflation. This cost can then be compared with the benefits of lower inflation.

¹⁰ Douglas Staiger, James H. Stock, and Mark W. Watson, "How Precise Are Estimates of the Natural Rate of Unemployment?" in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997).

Much research has used the available data to examine the Phillips curve quantitatively. The results of these studies are often summarized in a number called the **sacrifice ratio**, the percentage of a year's real GDP that must be forgone to reduce inflation by 1 percentage point. Although estimates of the sacrifice ratio vary substantially, a typical estimate is about 5: for every percentage point that inflation is to fall, 5 percent of one year's GDP must be sacrificed.¹¹

We can also express the sacrifice ratio in terms of unemployment. Okun's law says that a change of 1 percentage point in the unemployment rate translates into a change of 2 percentage points in GDP. Therefore, reducing inflation by 1 percentage point requires about 2.5 percentage points of cyclical unemployment.

We can use the sacrifice ratio to estimate by how much and for how long unemployment must rise to reduce inflation. If reducing inflation by 1 percentage point requires a sacrifice of 5 percent of a year's GDP, reducing inflation by 4 percentage points requires a sacrifice of 20 percent of a year's GDP. Equivalently, this reduction in inflation requires a sacrifice of 10 percentage points of cyclical unemployment.

This disinflation could take various forms, each totaling the same sacrifice of 20 percent of a year's GDP. For example, a rapid disinflation would lower output by 10 percent for 2 years: this is sometimes called the *cold-turkey* solution to inflation. A moderate disinflation would lower output by 5 percent for 4 years. An even more gradual disinflation would depress output by 2 percent for a decade.

Rational Expectations and the Possibility of Painless Disinflation

Because the expectation of inflation influences the short-run tradeoff between inflation and unemployment, it is crucial to understand how people form expectations. So far, we have been assuming that expected inflation depends on recently observed inflation. Although this assumption of adaptive expectations is plausible, it is probably too simple to apply in all circumstances.

An alternative approach is to assume that people have **rational expectations**. That is, we might assume that people optimally use all the available information, including information about current government policies, to forecast the future. Because monetary and fiscal policies influence inflation, expected inflation should also depend on the monetary and fiscal policies in effect. According to the theory of rational expectations, a change in monetary or fiscal policy will change expectations, and an evaluation of any policy change must incorporate this effect on expectations. If people do form their expectations rationally, then inflation may have less inertia than it first appears.

¹¹ Arthur M. Okun, "Efficient Disinflationary Policies," *American Economic Review* 68 (May 1978): 348–352; and Robert J. Gordon and Stephen R. King, "The Output Cost of Disinflation in Traditional and Vector Autoregressive Models," *Brookings Papers on Economic Activity* (1982:1): 205–245.

Here is how Thomas Sargent, a prominent advocate of rational expectations, describes its implications for the Phillips curve:

An alternative “rational expectations” view denies that there is any inherent momentum to the present process of inflation. This view maintains that firms and workers have now come to expect high rates of inflation in the future and that they strike inflationary bargains in light of these expectations. However, it is held that people expect high rates of inflation in the future precisely because the government’s current and prospective monetary and fiscal policies warrant those expectations. . . . Thus inflation only seems to have a momentum of its own; it is actually the long-term government policy of persistently running large deficits and creating money at high rates which imparts the momentum to the inflation rate. An implication of this view is that inflation can be stopped much more quickly than advocates of the “momentum” view have indicated and that their estimates of the length of time and the costs of stopping inflation in terms of foregone output are erroneous. . . . [Stopping inflation] would require a change in the policy regime: there must be an abrupt change in the continuing government policy, or strategy, for setting deficits now and in the future that is sufficiently binding as to be widely believed. . . . How costly such a move would be in terms of foregone output and how long it would be in taking effect would depend partly on how resolute and evident the government’s commitment was.¹²

Thus, advocates of rational expectations argue that the short-run Phillips curve does not accurately represent the options that policymakers have available. They believe that if policymakers are credibly committed to reducing inflation, rational people will understand the commitment and will quickly lower their expectations of inflation. Inflation can then come down without a rise in unemployment and fall in output. According to the theory of rational expectations, traditional estimates of the sacrifice ratio are not useful for evaluating the impact of alternative policies. Under a credible policy, the costs of reducing inflation may be much lower than estimates of the sacrifice ratio suggest.

In the most extreme case, one can imagine reducing the rate of inflation without causing any recession at all. A painless disinflation has two requirements. First, the plan to reduce inflation must be announced before the workers and firms who set wages and prices have formed their expectations. Second, the workers and firms must believe the announcement; otherwise, they will not reduce their expectations of inflation. If both requirements are met, the announcement will immediately shift the short-run tradeoff between inflation and unemployment downward, permitting a lower rate of inflation without higher unemployment.

Although the rational-expectations approach remains controversial, almost all economists agree that expectations of inflation influence the short-run tradeoff between inflation and unemployment. The credibility of a policy to reduce inflation is therefore one determinant of how costly the policy will be. Unfortunately, it is often difficult to predict whether the public will view the announcement of a new policy as credible. The central role of expectations makes forecasting the results of alternative policies far more difficult.

¹² Thomas J. Sargent, “The Ends of Four Big Inflation,” in Robert E. Hall, ed., *Inflation: Causes and Effects* (Chicago: University of Chicago Press, 1982).

CASE STUDY

The Sacrifice Ratio in Practice

The Phillips curve with adaptive expectations implies that reducing inflation requires a period of high unemployment and low output. By contrast, the rational-expectations approach suggests that reducing inflation can be much less costly. What happens during actual disinflations?

Consider the U.S. disinflation in the early 1980s. This decade began with some of the highest rates of inflation in U.S. history. Yet because of the tight monetary policies the Fed pursued under Chairman Paul Volcker, the rate of inflation fell substantially in the first few years of the decade. This episode provides a natural experiment with which to estimate how much output is lost during the process of disinflation.

The first question is, how much did inflation fall? As measured by the GDP deflator, inflation reached a peak of 9.7 percent in 1981. It is natural to end the episode in 1985 because oil prices plunged in 1986—a large, beneficial supply shock unrelated to Fed policy. In 1985, inflation was 3.0 percent, so we can estimate that the Fed engineered a reduction in inflation of 6.7 percentage points over four years.

The second question is, how much output was lost during this period? Table 13-1 shows the unemployment rate from 1982 to 1985. Assuming that the natural rate of unemployment was 6 percent, we can compute the amount of cyclical unemployment in each year. In total over this period, there were 9.5 percentage points of cyclical unemployment. Okun's law says that 1 percentage point of unemployment translates into 2 percentage points of GDP. Therefore, 19.0 percentage points of annual GDP were lost during the disinflation.

table 13-1

Unemployment During the Volcker Disinflation

Year	Unemployment Rate, u	Natural Rate, u^n	Cyclical Unemployment, $u - u^n$
1982	9.5%	6.0%	3.5%
1983	9.5	6.0	3.5
1984	7.4	6.0	1.4
1985	7.1	6.0	1.1
			Total 9.5%

Now we can compute the sacrifice ratio for this episode. We know that 19.0 percentage points of GDP were lost and that inflation fell by 6.7 percentage points. Hence, $19.0/6.7$, or 2.8, percentage points of GDP were lost for each percentage-point reduction in inflation. The estimate of the sacrifice ratio from the Volcker disinflation is 2.8.

This estimate of the sacrifice ratio is smaller than the estimates made before Volcker was appointed Fed chairman. In other words, Volcker reduced inflation at a smaller cost than many economists had predicted. One explanation is that Volcker's tough stand was credible enough to influence expectations of inflation directly. Yet the change in expectations was not large enough to make the disinflation painless: in 1982 unemployment reached its highest level since the Great Depression.

Although the Volcker disinflation is only one historical episode, this kind of analysis can be applied to other disinflations. A recent study documented the results of 65 disinflations in 19 countries. In almost all cases, the reduction in inflation came at the cost of temporarily lower output. Yet the size of the output loss varied from episode to episode. Rapid disinflations usually had smaller sacrifice ratios than slower ones. That is, in contrast to what the Phillips curve with adaptive expectations suggests, a cold-turkey approach appears less costly than a gradual one. Moreover, countries with more flexible wage-setting institutions, such as shorter labor contracts, had smaller sacrifice ratios. These findings indicate that reducing inflation always has some cost, but that policies and institutions can affect its magnitude.¹³

Hysteresis and the Challenge to the Natural-Rate Hypothesis

Our discussion of the cost of disinflation—and indeed our entire discussion of economic fluctuations in the past four chapters—has been based on an assumption called the **natural-rate hypothesis**. This hypothesis is summarized in the following statement:

Fluctuations in aggregate demand affect output and employment only in the short run. In the long run, the economy returns to the levels of output, employment, and unemployment described by the classical model.

The natural-rate hypothesis allows macroeconomists to study separately short-run and long-run developments in the economy. It is one expression of the classical dichotomy.

Recently, some economists have challenged the natural-rate hypothesis by suggesting that aggregate demand may affect output and employment even in the long run. They have pointed out a number of mechanisms through which recessions might leave permanent scars on the economy by altering the natural rate of unemployment. **Hysteresis** is the term used to describe the long-lasting influence of history on the natural rate.

A recession can have permanent effects if it changes the people who become unemployed. For instance, workers might lose valuable job skills when unemployed, lowering their ability to find a job even after the recession ends.

¹³ Laurence Ball, "What Determines the Sacrifice Ratio?" in N. Gregory Mankiw, ed., *Monetary Policy* (Chicago: University of Chicago Press, 1994).

Alternatively, a long period of unemployment may change an individual's attitude toward work and reduce his desire to find employment. In either case, the recession permanently inhibits the process of job search and raises the amount of frictional unemployment.

Another way in which a recession can permanently affect the economy is by changing the process that determines wages. Those who become unemployed may lose their influence on the wage-setting process. Unemployed workers may lose their status as union members, for example. More generally, some of the *insiders* in the wage-setting process become *outsiders*. If the smaller group of insiders cares more about high real wages and less about high employment, then the recession may permanently push real wages further above the equilibrium level and raise the amount of structural unemployment.

Hysteresis remains a controversial theory. Some economists believe the theory helps explain persistently high unemployment in Europe, because the rise in European unemployment starting in the early 1980s coincided with disinflation but continued after inflation stabilized. Moreover, the increase in unemployment tended to be larger for those countries that experienced the greatest reductions in inflations, such as Ireland, Italy, and Spain. Yet there is still no consensus whether the hysteresis phenomenon is significant, or why it might be more pronounced in some countries than in others. (Other explanations of high European unemployment, discussed in Chapter 6, give little role to the disinflation.) If it is true, however, the theory is important, because hysteresis greatly increases the cost of recessions. Put another way, hysteresis raises the sacrifice ratio, because output is lost even after the period of disinflation is over.¹⁴

13-3 Conclusion

We began this chapter by discussing three models of aggregate supply, each of which focuses on a different reason why the short-run aggregate supply curve is upward sloping. The three models have similar predictions for the aggregate economy, and all of them yield a short-run tradeoff between inflation and unemployment. A convenient way to express and analyze that tradeoff is with the Phillips curve equation, according to which inflation depends on expected inflation, cyclical unemployment, and supply shocks.

Keep in mind that not all economists endorse all the ideas discussed here. There is widespread disagreement, for instance, about the practical importance of rational expectations and the relevance of hysteresis. If you find it difficult to fit all the pieces together, you are not alone. The study of aggregate supply remains one of the most unsettled—and therefore one of the most exciting—research areas in macroeconomics.

¹⁴ Olivier J. Blanchard and Lawrence H. Summers, “Beyond the Natural Rate Hypothesis,” *American Economic Review* 78 (May 1988): 182–187; Laurence Ball, “Disinflation and the NAIRU,” in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997): 167–185.

Summary

1. The three theories of aggregate supply—the sticky-wage, imperfect-information, and sticky-price models—attribute deviations of output and employment from the natural rate to various market imperfections. According to all three theories, output rises above the natural rate when the price level exceeds the expected price level, and output falls below the natural rate when the price level is less than the expected price level.
2. Economists often express aggregate supply in a relationship called the Phillips curve. The Phillips curve says that inflation depends on expected inflation, the deviation of unemployment from its natural rate, and supply shocks. According to the Phillips curve, policymakers who control aggregate demand face a short-run tradeoff between inflation and unemployment.
3. If expected inflation depends on recently observed inflation, then inflation has inertia, which means that reducing inflation requires either a beneficial supply shock or a period of high unemployment and reduced output. If people have rational expectations, however, then a credible announcement of a change in policy might be able to influence expectations directly and, therefore, reduce inflation without causing a recession.
4. Most economists accept the natural-rate hypothesis, according to which fluctuations in aggregate demand have only short-run effects on output and unemployment. Yet some economists have suggested ways in which recessions can leave permanent scars on the economy by raising the natural rate of unemployment.

KEY CONCEPTS

Sticky-wage model	Adaptive expectations	Rational expectations
Imperfect-information model	Demand-pull inflation	Natural-rate hypothesis
Sticky-price model	Cost-push inflation	Hysteresis
Phillips curve	Sacrifice ratio	

QUESTIONS FOR REVIEW

1. Explain the three theories of aggregate supply. On what market imperfection does each theory rely? What do the theories have in common?
2. How is the Phillips curve related to aggregate supply?
3. Why might inflation be inertial?
4. Explain the differences between demand-pull inflation and cost-push inflation.
5. Under what circumstances might it be possible to reduce inflation without causing a recession?
6. Explain two ways in which a recession might raise the natural rate of unemployment.

PROBLEMS AND APPLICATIONS

1. Consider the following changes in the sticky-wage model.
 - a. Suppose that labor contracts specify that the nominal wage be fully indexed for inflation. That is, the nominal wage is to be adjusted to fully compensate for changes in the consumer price index. How does full indexation alter the aggregate supply curve in this model?
 - b. Suppose now that indexation is only partial. That is, for every increase in the CPI, the nominal wage rises, but by a smaller percentage. How does partial indexation alter the aggregate supply curve in this model?
2. In the sticky-price model, describe the aggregate supply curve in the following special cases. How do these cases compare to the short-run aggregate supply curve we discussed in Chapter 9?
 - a. No firms have flexible prices ($s = 1$).
 - b. The desired price does not depend on aggregate output ($a = 0$).
3. Suppose that an economy has the Phillips curve

$$\pi = \pi_{-1} - 0.5(u - 0.06).$$
 - a. What is the natural rate of unemployment?
 - b. Graph the short-run and long-run relationships between inflation and unemployment.
 - c. How much cyclical unemployment is necessary to reduce inflation by 5 percentage points? Using Okun's law, compute the sacrifice ratio.
 - d. Inflation is running at 10 percent. The Fed wants to reduce it to 5 percent. Give two scenarios that will achieve that goal.
4. According to the rational-expectations approach, if everyone believes that policymakers are committed to reducing inflation, the cost of reducing inflation—the sacrifice ratio—will be lower than if the public is skeptical about the policymakers' intentions. Why might this be true? How might credibility be achieved?
5. Assume that people have rational expectations and that the economy is described by the sticky-wage or sticky-price model. Explain why each of the following propositions is true:
 - a. Only unanticipated changes in the money supply affect real GDP. Changes in the money supply that were anticipated when wages and prices were set do not have any real effects.
 - b. If the Fed chooses the money supply at the same time as people are setting wages and prices, so that everyone has the same information about the state of the economy, then monetary policy cannot be used systematically to stabilize output. Hence, a policy of keeping the money supply constant will have the same real effects as a policy of adjusting the money supply in response to the state of the economy. (This is called the *policy irrelevance proposition*.)
 - c. If the Fed sets the money supply well after people have set wages and prices, so the Fed has collected more information about the state of the economy, then monetary policy can be used systematically to stabilize output.
6. Suppose that an economy has the Phillips curve

$$\pi = \pi_{-1} - 0.5(u - u^n),$$
 and that the natural rate of unemployment is given by an average of the past two years' unemployment:

$$u^n = 0.5(u_{-1} + u_{-2}).$$
 - a. Why might the natural rate of unemployment depend on recent unemployment (as is assumed in the preceding equation)?
 - b. Suppose that the Fed follows a policy to reduce permanently the inflation rate by 1 percentage point. What effect will that policy have on the unemployment rate over time?
 - c. What is the sacrifice ratio in this economy? Explain.
 - d. What do these equations imply about the short-run and long-run tradeoffs between inflation and unemployment?
7. Some economists believe that taxes have an important effect on labor supply. They argue that higher taxes cause people to want to work less and that lower taxes cause them to want to work more. Consider how this effect alters the macroeconomic analysis of tax changes.

- a. If this view is correct, how does a tax cut affect the natural rate of output?
 - b. How does a tax cut affect the aggregate demand curve? The long-run aggregate supply curve? The short-run aggregate supply curve?
 - c. What is the short-run impact of a tax cut on output and the price level? How does your answer differ from the case without the labor-supply effect?
 - d. What is the long-run impact of a tax cut on output and the price level? How does your answer differ from the case without the labor-supply effect?
8. Economist Alan Blinder, whom Bill Clinton appointed to be Vice Chairman of the Federal Reserve, once wrote the following:

The costs that attend the low and moderate inflation rates experienced in the United States and in other

industrial countries appear to be quite modest—more like a bad cold than a cancer on society. . . . As rational individuals, we do not volunteer for a lobotomy to cure a head cold. Yet, as a collectivity, we routinely prescribe the economic equivalent of lobotomy (high unemployment) as a cure for the inflationary cold.¹⁵

What do you think Blinder meant by this? What are the policy implications of the viewpoint Blinder is advocating? Do you agree? Why or why not?

9. Go to the Web site of the Bureau of Labor Statistics (www.bls.gov). For each of the past five years, find the inflation rate as measured by the consumer price index (all items) and as measured by the CPI excluding food and energy. Compare these two measures of inflation. Why might they be different? What might the difference tell you about shifts in the aggregate supply curve and in the short-run Phillips curve?

¹⁵ Alan Blinder, *Hard Heads, Soft Hearts: Tough-Minded Economics for a Just Society* (Reading, Mass.: Addison-Wesley, 1987), 51.

APPENDIX

A Big, Comprehensive Model

In the previous chapters, we have seen many models of how the economy works. When learning these models, it can be hard to see how they are related. Now that we have finished developing the model of aggregate demand and aggregate supply, this is a good time to look back at what we have learned. This appendix sketches a large model that incorporates much of the theory we have already seen, including the classical theory presented in Part II and the business cycle theory presented in Part IV. The notation and equations should be familiar from previous chapters.

The model has seven equations:

$Y = C(Y - T) + I(r) + G + NX(\epsilon)$	IS: Goods Market Equilibrium
$M/P = L(i, Y)$	LM: Money Market Equilibrium
$NX(\epsilon) = CF(r - r^*)$	Foreign Exchange Market Equilibrium
$i = r + \pi^e$	Relationship Between Real and Nominal Interest Rates
$\epsilon = eP/P^*$	Relationship Between Real and Nominal Exchange Rates
$Y = \bar{Y} + \alpha(P - P^e)$	Aggregate Supply
$\bar{Y} = F(\bar{K}, \bar{L})$	Natural Rate of Output

These seven equations determine the equilibrium values of seven endogenous variables: output Y , the natural rate of output $\bar{Y}$, the real interest rate r , the nominal interest rate i , the real exchange rate ϵ , the nominal exchange rate e , and the price level P .

There are many exogenous variables that influence these endogenous variables. They include the money supply M , government purchases G , taxes T , the capital stock K , the labor force L , the world price level P^* , and the world real interest rate r^* . In addition, there are two expectational variables: the expectation of future inflation π^e and the expectation of the current price level formed in the past P^e . As written, the model takes these expectational variables as exogenous, although additional equations could be added to make them endogenous.

Although mathematical techniques are available to analyze this seven-equation model, they are beyond the scope of this book. But this large model is still useful, because we can use it to see how the smaller models we have examined are related to one another. In particular, *many of the models we have been studying are special cases of this large model*. Let's consider six special cases.

Special Case 1: The Classical Closed Economy Suppose that $P^e = P$, $L(i, Y) = (1/V)Y$, and $CF(r - r^*) = 0$. In words, this means that expectations of the price

level adjust so that expectations are correct, that money demand is proportional to income, and that there are no international capital flows. In this case, output is always at its natural rate, the real interest rate adjusts to equilibrate the goods market, the price level moves parallel with the money supply, and the nominal interest rate adjusts one-for-one with expected inflation. This special case corresponds to the economy analyzed in Chapters 3 and 4.

Special Case 2: The Classical Small Open Economy Suppose that $P^e = P$, $L(i, Y) = (1/V)Y$, and $CF(r - r^*)$ is infinitely elastic. Now we are examining the special case when international capital flows respond greatly to any differences between the domestic and world interest rates. This means that $r = r^*$ and that the trade balance NX equals the difference between saving and investment at the world interest rate. This special case corresponds to the economy analyzed in Chapter 5.

Special Case 3: The Basic Model of Aggregate Demand and Aggregate Supply Suppose that α is infinite and $L(i, Y) = (1/V)Y$. In this case, the short-run aggregate supply curve is horizontal, and the aggregate demand curve is determined only by the quantity equation. This special case corresponds to the economy analyzed in Chapter 9.

Special Case 4: The IS-LM Model Suppose that α is infinite and $CF(r - r^*) = 0$. In this case, the short-run aggregate supply curve is horizontal, and there are no international capital flows. For any given level of expected inflation π^e , the level of income and interest rate must adjust to equilibrate the goods market and the money market. This special case corresponds to the economy analyzed in Chapter 10 and 11.

Special Case 5: The Mundell-Fleming Model with a Floating Exchange Rate Suppose that α is infinite and $CF(r - r^*)$ is infinitely elastic. In this case, the short-run aggregate supply curve is horizontal, and international capital flows are so great as to ensure that $r = r^*$. The exchange rate floats freely to reach its equilibrium level. This special case corresponds to the first economy analyzed in Chapter 12.

Special Case 6: The Mundell-Fleming Model with a Fixed Exchange Rate Suppose that α is infinite, $CF(r - r^*)$ is infinitely elastic, and e is fixed. In this case, the short-run aggregate supply curve is horizontal, huge international capital flows ensure that $r = r^*$, but the exchange rate is set by the central bank. The exchange rate is now an exogenous policy variable, but the money supply M is an endogenous variable that must adjust to ensure the exchange rate hits the fixed level. This special case corresponds to the second economy analyzed in Chapter 12.

You should now see the value in this big model. Even though the model is too large to be useful in developing an intuitive understanding of how the economy works, it shows that the different models we have been studying are closely related. Each model shows a different facet of the larger and more realistic model

presented here. In each chapter, we made some simplifying assumptions to make the big model smaller and easier to understand. When thinking about the real world, it is important to keep the simplifying assumptions in mind and to draw on the insights learned in each of the chapters.

MORE PROBLEMS AND APPLICATIONS

1. Let's consider some more special cases of this large model. Starting with the large model, what extra assumptions would you need to yield each of the following models:
 - a. The model of the classical large open economy in the appendix to Chapter 5.
 - b. The Keynesian cross in the first half of Chapter 10.
 - c. The *IS-LM* model for the large open economy in the appendix to Chapter 12.

part v

Macroeconomic Policy Debates

Stabilization Policy

The Federal Reserve's job is to take away the punch bowl just as the party gets going.

— William McChesney Martin

What we need is not a skilled monetary driver of the economic vehicle continuously turning the steering wheel to adjust to the unexpected irregularities of the route, but some means of keeping the monetary passenger who is in the back seat as ballast from occasionally leaning over and giving the steering wheel a jerk that threatens to send the car off the road.

— Milton Friedman

How should government policymakers respond to the business cycle? The two quotations above—the first from a former chairman of the Federal Reserve, the second from a prominent critic of the Fed—show the diversity of opinion over how this question is best answered.

Some economists, such as William McChesney Martin, view the economy as inherently unstable. They argue that the economy experiences frequent shocks to aggregate demand and aggregate supply. Unless policymakers use monetary and fiscal policy to stabilize the economy, these shocks will lead to unnecessary and inefficient fluctuations in output, unemployment, and inflation. According to the popular saying, macroeconomic policy should “lean against the wind,” stimulating the economy when it is depressed and slowing the economy when it is overheated.

Other economists, such as Milton Friedman, view the economy as naturally stable. They blame bad economic policies for the large and inefficient fluctuations we have sometimes experienced. They argue that economic policy should not try to “fine-tune” the economy. Instead, economic policymakers should admit their limited abilities and be satisfied if they do no harm.

This debate has persisted for decades, with numerous protagonists advancing various arguments for their positions. The fundamental issue is how policymakers should use the theory of short-run economic fluctuations developed in the

preceding chapters. In this chapter we ask two questions that arise in this debate. First, should monetary and fiscal policy take an active role in trying to stabilize the economy, or should policy remain passive? Second, should policymakers be free to use their discretion in responding to changing economic conditions, or should they be committed to following a fixed policy rule?

14-1 Should Policy Be Active or Passive?

Policymakers in the federal government view economic stabilization as one of their primary responsibilities. The analysis of macroeconomic policy is a regular duty of the Council of Economic Advisers, the Congressional Budget Office, the Federal Reserve, and other government agencies. When Congress or the president is considering a major change in fiscal policy, or when the Federal Reserve is considering a major change in monetary policy, foremost in the discussion are how the change will influence inflation and unemployment and whether aggregate demand needs to be stimulated or restrained.

Although the government has long conducted monetary and fiscal policy, the view that it should use these policy instruments to try to stabilize the economy is more recent. The Employment Act of 1946 was a key piece of legislation in which the government first held itself accountable for macroeconomic performance. The act states that “it is the continuing policy and responsibility of the Federal Government to . . . promote full employment and production.” This law was written when the memory of the Great Depression was still fresh. The lawmakers who wrote it believed, as many economists do, that in the absence of an active government role in the economy, events such as the Great Depression could occur regularly.

To many economists the case for active government policy is clear and simple. Recessions are periods of high unemployment, low incomes, and increased economic hardship. The model of aggregate demand and aggregate supply shows how shocks to the economy can cause recessions. It also shows how monetary and fiscal policy can prevent recessions by responding to these shocks. These economists consider it wasteful not to use these policy instruments to stabilize the economy.

Other economists are critical of the government’s attempts to stabilize the economy. These critics argue that the government should take a hands-off approach to macroeconomic policy. At first, this view might seem surprising. If our model shows how to prevent or reduce the severity of recessions, why do these critics want the government to refrain from using monetary and fiscal policy for economic stabilization? To find out, let’s consider some of their arguments.

Lags in the Implementation and Effects of Policies

Economic stabilization would be easy if the effects of policy were immediate. Making policy would be like driving a car: policymakers would simply adjust their instruments to keep the economy on the desired path.

Making economic policy, however, is less like driving a car than it is like piloting a large ship. A car changes direction almost immediately after the steering wheel is turned. By contrast, a ship changes course long after the pilot adjusts the rudder, and once the ship starts to turn, it continues turning long after the rudder is set back to normal. A novice pilot is likely to oversteer and, after noticing the mistake, overreact by steering too much in the opposite direction. The ship's path could become unstable, as the novice responds to previous mistakes by making larger and larger corrections.

Like a ship's pilot, economic policymakers face the problem of long lags. Indeed, the problem for policymakers is even more difficult, because the lengths of the lags are hard to predict. These long and variable lags greatly complicate the conduct of monetary and fiscal policy.

Economists distinguish between two lags in the conduct of stabilization policy: the inside lag and the outside lag. The **inside lag** is the time between a shock to the economy and the policy action responding to that shock. This lag arises because it takes time for policymakers first to recognize that a shock has occurred and then to put appropriate policies into effect. The **outside lag** is the time between a policy action and its influence on the economy. This lag arises because policies do not immediately influence spending, income, and employment.

A long inside lag is a central problem with using fiscal policy for economic stabilization. This is especially true in the United States, where changes in spending or taxes require the approval of the president and both houses of Congress. The slow and cumbersome legislative process often leads to delays, which make fiscal policy an imprecise tool for stabilizing the economy. This inside lag is shorter in countries with parliamentary systems, such as the United Kingdom, because there the party in power can often enact policy changes more rapidly.

Monetary policy has a much shorter inside lag than fiscal policy, because a central bank can decide on and implement a policy change in less than a day, but monetary policy has a substantial outside lag. Monetary policy works by changing the money supply and thereby interest rates, which in turn influence investment. But many firms make investment plans far in advance. Therefore, a change in monetary policy is thought not to affect economic activity until about six months after it is made.

The long and variable lags associated with monetary and fiscal policy certainly make stabilizing the economy more difficult. Advocates of passive policy argue that, because of these lags, successful stabilization policy is almost impossible. Indeed, attempts to stabilize the economy can be destabilizing. Suppose that the economy's condition changes between the beginning of a policy action and its impact on the economy. In this case, active policy may end up stimulating the economy when it is overheated or depressing the economy when it is cooling off. Advocates of active policy admit that such lags do require policymakers to be cautious. But, they argue, these lags do not necessarily mean that policy should be completely passive, especially in the face of a severe and protracted economic downturn.

Some policies, called **automatic stabilizers**, are designed to reduce the lags associated with stabilization policy. Automatic stabilizers are policies that stimulate or depress the economy when necessary without any deliberate policy change. For

example, the system of income taxes automatically reduces taxes when the economy goes into a recession, without any change in the tax laws, because individuals and corporations pay less tax when their incomes fall. Similarly, the unemployment-insurance and welfare systems automatically raise transfer payments when the economy moves into a recession, because more people apply for benefits. One can view these automatic stabilizers as a type of fiscal policy without any inside lag.

The Difficult Job of Economic Forecasting

Because policy influences the economy only after a long lag, successful stabilization policy requires the ability to predict accurately future economic conditions. If we cannot predict whether the economy will be in a boom or a recession in six months or a year, we cannot evaluate whether monetary and fiscal policy should now be trying to expand or contract aggregate demand. Unfortunately, economic developments are often unpredictable, at least given our current understanding of the economy.

One way forecasters try to look ahead is with **leading indicators**. A leading indicator is a data series that fluctuates in advance of the economy. A large fall in a leading indicator signals that a recession is more likely.

Another way forecasters look ahead is with macro-econometric models, which have been developed both by government agencies and by private firms for forecasting and policy analysis. As we discussed in Chapter 11,



Drawing by Dana Fradon; © 1988 The New Yorker Magazine, Inc.

“It’s true, Caesar. Rome is declining, but I expect it to pick up in the next quarter.”

FYI

What’s in the Index of Leading Economic Indicators?

Each month the Conference Board, a private economics research group, announces the *index of leading economic indicators*. This index is made up from 10 data series that are often used to forecast changes in economic activity about six to nine months ahead. Here is a list of the series. Can you explain why each of these might help predict changes in real GDP?

1. Average workweek of production workers in manufacturing.
2. Average initial weekly claims for unemployment insurance. This series is inverted in computing the index, so that a decrease in the series raises the index.
3. New orders for consumer goods and materials, adjusted for inflation.
4. Vendor performance. This is a measure of the number of companies receiving slower deliveries from suppliers.
5. New orders, nondefense capital goods.
6. New building permits issued.
7. Index of stock prices.
8. Money supply ($M2$), adjusted for inflation.
9. Interest rate spread: the yield spread between 10-year Treasury notes and 3-month Treasury bills.
10. Index of consumer expectations.


these large-scale computer models are made up of many equations, each representing a part of the economy. After making assumptions about the path of the exogenous variables, such as monetary policy, fiscal policy, and oil prices, these models yield predictions about unemployment, inflation, and other endogenous variables. Keep in mind, however, that the validity of these predictions is only as good as the model and the forecasters' assumptions about the exogenous variables.

CASE STUDY

Mistakes in Forecasting

“Light showers, bright intervals, and moderate winds.” This was the forecast offered by the renowned British national weather service on October 14, 1987. The next day Britain was hit by the worst storm in more than two centuries.

figure 14-1



Forecasting the Recession of 1982 The red line shows the actual unemployment rate from the first quarter of 1980 to the first quarter of 1986. The blue lines show the unemployment rate predicted at six points in time: the second quarter of 1981, the fourth quarter of 1981, the second quarter of 1982, and so on. For each forecast, the symbols mark the current unemployment rate and the forecast for the subsequent five quarters. Notice that the forecasters failed to predict both the rapid rise in the unemployment rate and the subsequent rapid decline.

Source: The unemployment rate is from the Department of Commerce. The predicted unemployment rate is the median forecast of about 20 forecasters surveyed by the American Statistical Association and the National Bureau of Economic Research.

Like weather forecasts, economic forecasts are a crucial input to private and public decisionmaking. Business executives rely on economic forecasts when deciding how much to produce and how much to invest in plant and equipment. Government policymakers also rely on them when developing economic policies. Yet also like weather forecasts, economic forecasts are far from precise.

The most severe economic downturn in U.S. history, the Great Depression of the 1930s, caught economic forecasters completely by surprise. Even after the stock market crash of 1929, they remained confident that the economy would not suffer a substantial setback. In late 1931, when the economy was clearly in bad shape, the eminent economist Irving Fisher predicted that it would recover quickly. Subsequent events showed that these forecasts were much too optimistic.¹

Figure 14-1 shows how economic forecasters did during the recession of 1982, the most severe economic downturn in the United States since the Great Depression. This figure shows the actual unemployment rate (in red) and six attempts to predict it for the following five quarters (in blue). You can see that the forecasters did well predicting unemployment one quarter ahead. The more distant forecasts, however, were often inaccurate. For example, in the second quarter of 1981, forecasters were predicting little change in the unemployment rate over the next five quarters; yet only two quarters later unemployment began to rise sharply. The rise in unemployment to almost 11 percent in the fourth quarter of 1982 caught the forecasters by surprise. After the depth of the recession became apparent, the forecasters failed to predict how rapid the subsequent decline in unemployment would be.

These two episodes—the Great Depression and the recession of 1982—show that many of the most dramatic economic events are unpredictable. Although private and public decisionmakers have little choice but to rely on economic forecasts, they must always keep in mind that these forecasts come with a large margin of error.

Ignorance, Expectations, and the Lucas Critique

The prominent economist Robert Lucas once wrote, “As an advice-giving profession we are in way over our heads.” Even many of those who advise policymakers would agree with this assessment. Economics is a young science, and there is still much that we do not know. Economists cannot be completely confident when they assess the effects of alternative policies. This ignorance suggests that economists should be cautious when offering policy advice.

Although economists’ knowledge is limited about many topics, Lucas has emphasized the issue of how people form expectations of the future. Expectations play a crucial role in the economy because they influence all sorts of economic behavior. For instance, households decide how much to consume based on expectations

¹ Kathryn M. Dominguez, Ray C. Fair, and Matthew D. Shapiro, “Forecasting the Depression: Harvard Versus Yale,” *American Economic Review* 78 (September 1988): 595–612. This article shows how badly economic forecasters did during the Great Depression, and it argues that they could not have done any better with the modern forecasting techniques available today.

of future income, and firms decide how much to invest based on expectations of future profitability. These expectations depend on many things, including the economic policies being pursued by the government. Thus, when policymakers estimate the effect of any policy change, they need to know how people's expectations will respond to the policy change. Lucas has argued that traditional methods of policy evaluation—such as those that rely on standard macroeconomic models—do not adequately take into account this impact of policy on expectations. This criticism of traditional policy evaluation is known as the **Lucas critique**.²

An important example of the Lucas critique arises in the analysis of disinflation. As you may recall from Chapter 13, the cost of reducing inflation is often measured by the sacrifice ratio, which is the number of percentage points of GDP that must be forgone to reduce inflation by 1 percentage point. Because these estimates of the sacrifice ratio are often large, they have led some economists to argue that policymakers should learn to live with inflation, rather than incur the large cost of reducing it.

According to advocates of the rational-expectations approach, however, these estimates of the sacrifice ratio are unreliable because they are subject to the Lucas critique. Traditional estimates of the sacrifice ratio are based on adaptive expectations, that is, on the assumption that expected inflation depends on past inflation. Adaptive expectations may be a reasonable premise in some circumstances, but if the policymakers make a credible change in policy, workers and firms setting wages and prices will rationally respond by adjusting their expectations of inflation appropriately. This change in inflation expectations will quickly alter the short-run tradeoff between inflation and unemployment. As a result, reducing inflation can potentially be much less costly than is suggested by traditional estimates of the sacrifice ratio.

The Lucas critique leaves us with two lessons. The narrow lesson is that economists evaluating alternative policies need to consider how policy affects expectations and, thereby, behavior. The broad lesson is that policy evaluation is hard, so economists engaged in this task should be sure to show the requisite humility.

The Historical Record

In judging whether government policy should play an active or passive role in the economy, we must give some weight to the historical record. If the economy has experienced many large shocks to aggregate supply and aggregate demand, and if policy has successfully insulated the economy from these shocks, then the case for active policy should be clear. Conversely, if the economy has experienced few large shocks, and if the fluctuations we have observed can be traced to inept economic policy, then the case for passive policy should be clear. In other words, our view of stabilization policy should be influenced by whether policy has historically been stabilizing or destabilizing. For this reason, the debate over macroeconomic policy frequently turns into a debate over macroeconomic history.

Yet history does not settle the debate over stabilization policy. Disagreements over history arise because it is not easy to identify the sources of

² Robert E. Lucas, Jr., "Econometric Policy Evaluation: A Critique," *Carnegie Rochester Conference on Public Policy* 1 (Amsterdam: North-Holland, 1976), 19–46.

economic fluctuations. The historical record often permits more than one interpretation.

The Great Depression is a case in point. Economists' views on macroeconomic policy are often related to their views on the cause of the Depression. Some economists believe that a large contractionary shock to private spending caused the Depression. They assert that policymakers should have responded by stimulating aggregate demand. Other economists believe that the large fall in the money supply caused the Depression. They assert that the Depression would have been avoided if the Fed had been pursuing a passive monetary policy of increasing the money supply at a steady rate. Hence, depending on one's beliefs about its cause, the Great Depression can be viewed either as an example of why active monetary and fiscal policy is necessary or as an example of why it is dangerous.

CASE STUDY

Is the Stabilization of the Economy a Figment of the Data?

Keynes wrote *The General Theory* in the 1930s, and in the wake of the Keynesian revolution, governments around the world began to view economic stabilization as a primary responsibility. Some economists believe that the development of Keynesian theory has had a profound influence on the behavior of the economy. Comparing data from before World War I and after World War II, they find that real GDP and unemployment have become much more stable. This, some Keynesians claim, is the best argument for active stabilization policy: it has worked.

In a series of provocative and influential papers, economist Christina Romer has challenged this assessment of the historical record. She argues that the measured reduction in volatility reflects not an improvement in economic policy and performance but rather an improvement in the economic data. The older data are much less accurate than the newer data. Romer claims that the higher volatility of unemployment and real GDP reported for the period before World War I is largely a figment of the data.

Romer uses various techniques to make her case. One is to construct more accurate data for the earlier period. This task is difficult because data sources are not readily available. A second way is to construct *less* accurate data for the recent period—that is, data that are comparable to the older data and thus suffer from the same imperfections. After constructing new “bad” data, Romer finds that the recent period appears almost as volatile as the early period, suggesting that the volatility of the early period may be largely an artifact of data construction.

Romer's work is part of the continuing debate over whether macroeconomic policy has improved the performance of the economy. Although her work remains controversial, most economists now believe that the economy in the aftermath of the Keynesian revolution was only slightly more stable than it had been before.³

³ Christina D. Romer, “Spurious Volatility in Historical Unemployment Data,” *Journal of Political Economy* 94 (February 1986): 1–37; and Christina D. Romer, “Is the Stabilization of the Postwar Economy a Figment of the Data?” *American Economic Review* 76 (June 1986): 314–334.

CASE STUDY

The Remarkable Stability of the 1990s

Although economists who take a long historical view debate how much the economy has stabilized over time, there is less controversy about the more recent experience. Everyone agrees that the decade of the 1990s stands out as a period of remarkable stability for the U.S. economy.

Table 14-1 presents some statistics about economic performance for the last five decades of the twentieth century. The three variables are those highlighted in Chapter 2: growth in real GDP, inflation, and unemployment. For each variable, the table presents the average over each decade and the standard deviation. The standard deviation measures the volatility of a variable: the higher the standard deviation, the more volatile the variable is.

table 14-1

The U.S. Macroeconomic Experience, Decade by Decade

	1950s	1960s	1970s	1980s	1990s
Real GDP Growth					
Average	4.18	4.43	3.28	3.02	3.03
Standard deviation	3.89	2.13	2.80	2.68	1.56
Inflation					
Average	2.07	2.33	7.09	5.66	3.00
Standard deviation	2.44	1.48	2.72	3.53	1.12
Unemployment					
Average	4.51	4.78	6.22	7.27	5.76
Standard deviation	1.29	1.07	1.16	1.48	1.05

Note: Real GDP growth is the growth rate of inflation-adjusted gross domestic product from four quarters earlier. Inflation is the rate of change in the consumer price index over the previous 12 months. Unemployment is the monthly, seasonally adjusted percentage of the labor force without a job.

Source: Department of Commerce, Department of Labor, and author's calculations.

One striking result from this table is the low volatility of the 1990s. The averages for real growth, inflation, or unemployment are not unusual by historical standards, but the standard deviations of these variables are the smallest ever seen. Moreover, the changes are large. For instance, the standard deviation of inflation was 24 percent lower in the 1990s than it was in the 1960s—the second most stable decade.

What accounts for the stability of the 1990s? Part of the answer is luck. The U.S. economy did not have to deal with any large, adverse supply shocks, such as the oil-price shocks of the 1970s. Part of the answer is also good policy. Many economists give credit to Alan Greenspan, who was chairman of the Federal

Reserve throughout the 1990s. His decisions about interest rates and the money supply kept the economy on track, avoiding both deep recession and runaway inflation.

14-2 Should Policy Be Conducted by Rule or by Discretion?

A second topic of debate among economists is whether economic policy should be conducted by rule or by discretion. Policy is conducted by rule if policymakers announce in advance how policy will respond to various situations and commit themselves to following through on this announcement. Policy is conducted by discretion if policymakers are free to size up events as they occur and choose whatever policy seems appropriate at the time.

The debate over rules versus discretion is distinct from the debate over passive versus active policy. Policy can be conducted by rule and yet be either passive or active. For example, a passive policy rule might specify steady growth in the money supply of 3 percent per year. An active policy rule might specify that

$$\text{Money Growth} = 3\% + (\text{Unemployment Rate} - 6\%).$$

Under this rule, the money supply grows at 3 percent if the unemployment rate is 6 percent, but for every percentage point by which the unemployment rate exceeds 6 percent, money growth increases by an extra percentage point. This rule tries to stabilize the economy by raising money growth when the economy is in a recession.

We begin this section by discussing why policy might be improved by a commitment to a policy rule. We then examine several possible policy rules.

Distrust of Policymakers and the Political Process

Some economists believe that economic policy is too important to be left to the discretion of policymakers. Although this view is more political than economic, evaluating it is central to how we judge the role of economic policy. If politicians are incompetent or opportunistic, then we may not want to give them the discretion to use the powerful tools of monetary and fiscal policy.

Incompetence in economic policy arises for several reasons. Some economists view the political process as erratic, perhaps because it reflects the shifting power of special interest groups. In addition, macroeconomics is complicated, and politicians often do not have sufficient knowledge of it to make informed judgments. This ignorance allows charlatans to propose incorrect but superficially appealing solutions to complex problems. The political process often cannot weed out the advice of charlatans from that of competent economists.

Opportunism in economic policy arises when the objectives of policymakers conflict with the well-being of the public. Some economists fear that politicians

use macroeconomic policy to further their own electoral ends. If citizens vote on the basis of economic conditions prevailing at the time of the election, then politicians have an incentive to pursue policies that will make the economy look good during election years. A president might cause a recession soon after coming into office to lower inflation and then stimulate the economy as the next election approaches to lower unemployment; this would ensure that both inflation and unemployment are low on election day. Manipulation of the economy for electoral gain, called the **political business cycle**, has been the subject of extensive research by economists and political scientists.⁴

Distrust of the political process leads some economists to advocate placing economic policy outside the realm of politics. Some have proposed constitutional amendments, such as a balanced-budget amendment, that would tie the hands of legislators and insulate the economy from both incompetence and opportunism.

CASE STUDY

The Economy Under Republican and Democratic Presidents

How does the political party in power affect the economy? Researchers working at the boundary between economics and political science have been studying this question. One intriguing finding is that the two political parties in the United States appear to conduct systematically different macroeconomic policies.

Table 14-2 presents the growth in real GDP in each of the four years of the presidential terms since 1948. Notice that growth is usually low, and often negative, in the second year of Republican administrations. Six of the eight years in which real GDP fell are second or third years of Republican administrations. By contrast, the economy is usually booming in the second and third years of Democratic administrations.

One interpretation of this finding is that the two parties have different preferences regarding inflation and unemployment. That is, rather than viewing politicians as opportunistic, perhaps we should view them as merely partisan. Republicans seem to dislike inflation more than Democrats do. Therefore, Republicans pursue contractionary policies soon after coming into office and are willing to endure a recession to reduce inflation. Democrats pursue more expansionary policies to reduce unemployment and are willing to endure the higher inflation that results. Examining growth in the money supply shows that monetary policy is, in fact, less inflationary during Republican administrations. Thus, it seems that the two political parties pursue dramatically different policies and that the political process is one source of economic fluctuations.

Even if we accept this interpretation of the evidence, it is not clear whether it argues for or against fixed policy rules. On the one hand, a policy rule would

⁴ William Nordhaus, "The Political Business Cycle," *Review of Economic Studies* 42 (1975): 169–190; and Edward Tufte, *Political Control of the Economy* (Princeton, N. J.: Princeton University Press, 1978).

table 14-2

Real GDP Growth During Republican and Democratic Administrations

President	YEAR OF TERM			
	First	Second	Third	Fourth
Democratic Administrations				
Truman	-0.6	8.9	7.6	3.7
Kennedy/Johnson	2.3	6.0	4.3	5.8
Johnson	6.4	6.6	2.5	4.8
Carter	4.6	5.5	3.2	-0.2
Clinton I	2.7	4.0	2.7	3.6
Clinton II	4.4	4.4	4.2	5.0
Average	3.3	5.9	4.1	3.8
Republican Administrations				
Eisenhower I	4.6	-0.7	7.1	2.0
Eisenhower II	2.0	-1.0	7.2	2.5
Nixon	3.0	0.2	3.3	5.4
Nixon/Ford	5.8	-0.6	-0.4	5.6
Reagan I	2.5	-2.0	4.3	7.3
Reagan II	3.8	3.4	3.4	4.2
Bush (elder)	3.5	1.8	-0.5	3.0
Average	3.6	0.2	3.5	4.3

Source: Department of Commerce

insulate the economy from these political shocks. Under a fixed rule, the Fed would be unable to alter monetary policy in response to the changing political climate. The economy might be more stable, and long-run economic performance might be improved. On the other hand, a fixed policy rule would reduce the voice of the electorate in influencing macroeconomic policy.⁵

The Time Inconsistency of Discretionary Policy

If we assume that we can trust our policymakers, discretion at first glance appears superior to a fixed policy rule. Discretionary policy is, by its nature, flexible. As long as policymakers are intelligent and benevolent, there might appear to be little reason to deny them flexibility in responding to changing conditions.

Yet a case for rules over discretion arises from the problem of **time inconsistency** of policy. In some situations policymakers may want to announce in

⁵ Alberto Alesina, "Macroeconomics and Politics," *NBER Macroeconomics Annual* 3 (1988): 13–52.

advance the policy they will follow in order to influence the expectations of private decisionmakers. But later, after the private decisionmakers have acted on the basis of their expectations, these policymakers may be tempted to renege on their announcement. Understanding that policymakers may be inconsistent over time, private decisionmakers are led to distrust policy announcements. In this situation, to make their announcements credible, policymakers may want to make a commitment to a fixed policy rule.

Time inconsistency is illustrated most simply in a political rather than an economic example—specifically, public policy about negotiating with terrorists over the release of hostages. The announced policy of many nations is that they will not negotiate over hostages. Such an announcement is intended to deter terrorists: if there is nothing to be gained from kidnapping hostages, rational terrorists won't kidnap any. In other words, the purpose of the announcement is to influence the expectations of terrorists and thereby their behavior.

But, in fact, unless the policymakers are credibly committed to the policy, the announcement has little effect. Terrorists know that once hostages are taken, policymakers face an overwhelming temptation to make some concession to obtain the hostages' release. The only way to deter rational terrorists is to take away the discretion of policymakers and commit them to a rule of never negotiating. If policymakers were truly unable to make concessions, the incentive for terrorists to take hostages would be largely eliminated.

The same problem arises less dramatically in the conduct of monetary policy. Consider the dilemma of a Federal Reserve that cares about both inflation and unemployment. According to the Phillips curve, the tradeoff between inflation and unemployment depends on expected inflation. The Fed would prefer everyone to expect low inflation so that it will face a favorable tradeoff. To reduce expected inflation, the Fed might announce that low inflation is the paramount goal of monetary policy.

But an announcement of a policy of low inflation is by itself not credible. Once households and firms have formed their expectations of inflation and set wages and prices accordingly, the Fed has an incentive to renege on its announcement and implement expansionary monetary policy to reduce unemployment. People understand the Fed's incentive to renege and therefore do not believe the announcement in the first place. Just as a president facing a hostage crisis is sorely tempted to negotiate their release, a Federal Reserve with discretion is sorely tempted to inflate in order to reduce unemployment. And just as terrorists discount announced policies of never negotiating, households and firms discount announced policies of low inflation.

The surprising outcome of this analysis is that policymakers can sometimes better achieve their goals by having their discretion taken away from them. In the case of rational terrorists, fewer hostages will be taken and killed if policymakers are committed to following the seemingly harsh rule of refusing to negotiate for hostages' freedom. In the case of monetary policy, there will be lower inflation without higher unemployment if the Fed is committed to a policy of zero inflation. (This conclusion about monetary policy is modeled more explicitly in the appendix to this chapter.)

The time inconsistency of policy arises in many other contexts. Here are some examples:

- To encourage investment, the government announces that it will not tax income from capital. But after factories have been built, the government is tempted to renege on its promise to raise more tax revenue from them.
- To encourage research, the government announces that it will give a temporary monopoly to companies that discover new drugs. But after a drug has been discovered, the government is tempted to revoke the patent or to regulate the price to make the drug more affordable.
- To encourage good behavior, a parent announces that he or she will punish a child whenever the child breaks a rule. But after the child has misbehaved, the parent is tempted to forgive the transgression, because punishment is unpleasant for the parent as well as for the child.
- To encourage you to work hard, your professor announces that this course will end with an exam. But after you have studied and learned all the material, the professor is tempted to cancel the exam so that he or she won't have to grade it.

In each case, rational agents understand the incentive for the policymaker to renege, and this expectation affects their behavior. And in each case, the solution is to take away the policymaker's discretion with a credible commitment to a fixed policy rule.

CASE STUDY

Alexander Hamilton Versus Time Inconsistency

Time inconsistency has long been a problem associated with discretionary policy. In fact, it was one of the first problems that confronted Alexander Hamilton when President George Washington appointed him the first U.S. Secretary of the Treasury in 1789.

Hamilton faced the question of how to deal with the debts that the new nation had accumulated as it fought for its independence from Britain. When the revolutionary government incurred the debts, it promised to honor them when the war was over. But after the war, many Americans advocated defaulting on the debt because repaying the creditors would require taxation, which is always costly and unpopular.

Hamilton opposed the time-inconsistency policy of repudiating the debt. He knew that the nation would likely need to borrow again sometime in the future. In his *First Report on the Public Credit*, which he presented to Congress in 1790, he wrote

If the maintenance of public credit, then, be truly so important, the next inquiry which suggests itself is: By what means is it to be effected? The ready answer to which question is, by good faith; by a punctual performance of contracts. States,

like individuals, who observe their engagements are respected and trusted, while the reverse is the fate of those who pursue an opposite conduct.

Thus, Hamilton proposed that the nation make a commitment to the policy rule of honoring its debts.

The policy rule that Hamilton originally proposed has continued for more than two centuries. Today, unlike in Hamilton's time, when Congress debates spending priorities, no one seriously proposes defaulting on the public debt as a way to reduce taxes. In the case of public debt, everyone now agrees that the government should be committed to a fixed policy rule.

Rules for Monetary Policy

Even if we are convinced that policy rules are superior to discretion, the debate over macroeconomic policy is not over. If the Fed were to commit to a rule for monetary policy, what rule should it choose? Let's discuss briefly three policy rules that various economists advocate.

Some economists, called **monetarists**, advocate that the Fed keep the money supply growing at a steady rate. The quotation at the beginning of this chapter from Milton Friedman—the most famous monetarist—exemplifies this view of monetary policy. Monetarists believe that fluctuations in the money supply are responsible for most large fluctuations in the economy. They argue that slow and steady growth in the money supply would yield stable output, employment, and prices.

Although a monetarist policy rule might have prevented many of the economic fluctuations we have experienced historically, most economists believe that it is not the best possible policy rule. Steady growth in the money supply stabilizes aggregate demand only if the velocity of money is stable. But sometimes the economy experiences shocks, such as shifts in money demand, that cause velocity to be unstable. Most economists believe that a policy rule needs to allow the money supply to adjust to various shocks to the economy.

A second policy rule that economists widely advocate is nominal GDP targeting. Under this rule, the Fed announces a planned path for nominal GDP. If nominal GDP rises above the target, the Fed reduces money growth to dampen aggregate demand. If it falls below the target, the Fed raises money growth to stimulate aggregate demand. Because a nominal GDP target allows monetary policy to adjust to changes in the velocity of money, most economists believe it would lead to greater stability in output and prices than a monetarist policy rule.

A third policy rule that is often advocated is inflation targeting. Under this rule, the Fed would announce a target for the inflation rate (usually a low one) and then adjust the money supply when the actual inflation deviates from the target. Like nominal GDP targeting, inflation targeting insulates the economy from changes in the velocity of money. In addition, an inflation target has the political advantage that it is easy to explain to the public.

Notice that all these rules are expressed in terms of some nominal variable—the money supply, nominal GDP, or the price level. One can also imagine policy rules expressed in terms of real variables. For example, the Fed might try to target the unemployment rate at 5 percent. The problem with such a rule is that no one

knows exactly what the natural rate of unemployment is. If the Fed chose a target for the unemployment rate below the natural rate, the result would be accelerating inflation. Conversely, if the Fed chose a target for the unemployment rate above the natural rate, the result would be accelerating deflation. For this reason, economists rarely advocate rules for monetary policy expressed solely in terms of real variables, even though real variables such as unemployment and real GDP are the best measures of economic performance.

CASE STUDY

Inflation Targeting: Rule or Constrained Discretion?

Since the late 1980s, many of the world's central banks—including those of Australia, Canada, Finland, Israel, New Zealand, Spain, Sweden, and the United Kingdom—have adopted some form of an inflation target. Sometimes inflation targeting takes the form of a central bank announcing its policy intentions. Other times it takes the form of a national law that spells out the goals of monetary policy. For example, the Reserve Bank of New Zealand Act of 1989 told the central bank “to formulate and implement monetary policy directed to the economic objective of achieving and maintaining stability in the general level of prices.” The act conspicuously omitted any mention of any other competing objective, such as stability in output, employment, interest rates, or exchange rates. Although the U.S. Federal Reserve has not adopted inflation targeting, some members of Congress have proposed bills that would require the Fed to do so.

Should we interpret inflation targeting as a type of precommitment to a policy rule? Not completely. In all the countries that have adopted inflation targeting, central banks are left with a fair amount of discretion. Inflation targets are usually set as a range—an inflation rate of 1 to 3 percent, for instance—rather than a particular number. Thus, the central bank can choose where in the range it wants to be. In addition, the central banks are sometimes allowed to adjust their targets for inflation, at least temporarily, if some exogenous event (such as an easily identified supply shock) pushes inflation outside of the range that was previously announced.

In light of this flexibility, what is the purpose of inflation targeting? Although inflation targeting does leave the central bank with some discretion, the policy does constrain how this discretion is used. When a central bank is told to “do the right thing,” it is hard to hold the central bank accountable, because people can argue forever about what the right thing is in any specific circumstance. By contrast, when a central bank has announced an inflation target, the public can more easily judge whether the central bank is meeting that target. Thus, although inflation targeting does not tie the hands of the central bank, it does increase the transparency of monetary policy and, by doing so, makes central bankers more accountable for their actions.⁶

⁶ See Ben S. Bernanke and Frederic S. Mishkin, “Inflation Targeting: A New Framework for Monetary Policy?” *Journal of Economic Perspectives* 11 (Spring 1997): 97–116.

CASE STUDY

John Taylor's (and Alan Greenspan's?) Rule for Monetary Policy

If you wanted to set interest rates to achieve stable prices while avoiding large fluctuations in output and employment, how would you do it? This is exactly the question that Alan Greenspan and the other governors of the Federal Reserve must ask themselves every day. The short-term policy instrument that the Fed now sets is the federal funds rate—the short-term interest rate at which banks make loans to one another. Whenever the Federal Open Market Committee meets, it chooses a target for the federal funds rate. The Fed's bond traders are then told to conduct open-market operations in order to hit the desired target.

The hard part of the Fed's job is choosing the target for the federal funds rate. Two guidelines are clear. First, when inflation heats up, the federal funds rate should rise. An increase in the interest rate will mean a smaller money supply and, eventually, lower investment, lower output, higher unemployment, and reduced inflation. Second, when real economic activity slows—as reflected in real GDP or unemployment—the federal funds rate should fall. A decrease in the interest rate will mean a larger money supply and, eventually, higher investment, higher output, and lower unemployment.

The Fed needs to go beyond these general guidelines, however, and decide how much to respond to changes in inflation and real economic activity. To help it make this decision, economist John Taylor has proposed a simple rule for the federal funds rate:

$$\begin{aligned}\text{Nominal Federal Funds Rate} &= \text{Inflation} + 2.0 \\ &\quad + 0.5 (\text{Inflation} - 2.0) - 0.5 (\text{GDP Gap}).\end{aligned}$$


The GDP gap is the percentage shortfall of real GDP from an estimate of its natural rate.

Taylor's rule has the real federal funds rate—the nominal rate minus inflation—responding to inflation and the GDP gap. According to this rule, the real federal funds rate equals 2 percent when inflation is 2 percent and GDP is at its natural rate. For each percentage point by which inflation rises above 2 percent, the real federal funds rate rises by 0.5 percent. For each percentage point by which real GDP falls below its natural rate, the real federal funds rate falls by 0.5 percent. If GDP rises above its natural rate, so that the GDP gap is negative, the real federal funds rate rises accordingly.

One way to view the Taylor rule is as a complement to (rather than a substitute for) inflation targeting. As the previous case study discussed, inflation targeting offers a plan for the central bank in the medium run, but it does not constrain its month-to-month policy decisions. The Taylor rule may be a good short-run operating procedure for hitting a medium-run inflation target. According to the Taylor rule, monetary policy responds directly to inflation—as any inflation-targeting central bank must. But it also responds to the output gap, which can be viewed as a measure of inflationary pressures.

Taylor's rule for monetary policy is not only simple and reasonable, but it also resembles actual Fed behavior in recent years. Figure 14-2 shows the actual fed-

figure 14-2



The Federal Funds Rate: Actual and Suggested This figure shows the federal funds rate—the short-term interest rate at which banks make loans to each other. It also shows the federal funds rate suggested by John Taylor’s monetary rule. Notice that the two series move closely together.

Source: Federal Reserve Board, U.S. Department of Commerce, and author’s calculations. To implement the Taylor rule, the inflation rate is measured as the percentage change in the GDP deflator over the previous four quarters, and the GDP gap is measured as twice the deviation of the unemployment rate from 6 percent.

eral funds rate and the target rate as determined by Taylor’s proposed rule. Notice how closely together the two series move. John Taylor’s monetary rule may be more than an academic suggestion. It may be the rule that Alan Greenspan and his colleagues subconsciously follow.⁷

CASE STUDY

Central-Bank Independence

Suppose you were put in charge of writing the constitution and laws for a country. Would you give the president of the country authority over the policies of the central bank? Or would you allow the central bank to make decisions free from such political influence? In other words, assuming that monetary policy is made by discretion rather than by rule, who should exercise that discretion?

Countries vary greatly in how they choose to answer this question. In some countries, the central bank is a branch of the government; in others, the central

⁷ John B. Taylor, “The Inflation/Output Variability Tradeoff Revisited,” in *Goals, Guidelines, and Constraints Facing Monetary Policymakers* (Federal Reserve Bank of Boston, 1994).


bank is largely independent. In the United States, Fed governors are appointed by the president for 14-year terms, and they cannot be recalled if the president is unhappy with their decisions. This institutional structure gives the Fed a degree of independence similar to that of the Supreme Court.

Many researchers have investigated the effects of constitutional design on monetary policy. They have examined the laws of different countries to construct an index of central-bank independence. This index is based on various characteristics, such as the length of bankers' terms, the role of government officials on the bank board, and the frequency of contact between the government and the central bank. The researchers then examined the correlation between central-bank independence and macroeconomic performance.

The results of these studies are striking: more independent central banks are strongly associated with lower and more stable inflation. Figure 14-3 shows a scatterplot of central-bank independence and average inflation for the period 1955 to 1988. Countries that had an independent central bank, such as Germany, Switzerland, and the United States, tended to have low average inflation. Countries that had central banks with less independence, such as New Zealand and Spain, tended to have higher average inflation.

Researchers have also found there is no relationship between central-bank independence and real economic activity. In particular, central-bank independence is not correlated with average unemployment, the volatility of unem-

figure 14-3



Inflation and Central-Bank Independence This scatterplot presents the international experience with central-bank independence. The evidence shows that more-independent central banks tend to produce lower rates of inflation.

Source: Figure 1a, page 155, of Alberto Alesina and Lawrence H. Summers, "Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence," *Journal of Money, Credit, and Banking* 25 (May 1993): 151–162. Average inflation is for the period 1955–1988.

ployment, the average growth of real GDP, or the volatility of real GDP. Central-bank independence appears to offer countries a free lunch: it has the benefit of lower inflation without any apparent cost. This finding has led some countries, such as New Zealand, to rewrite their laws to give their central banks greater independence.⁸

14-3 Conclusion: Making Policy in an Uncertain World

In this chapter we have examined whether policy should take an active or passive role in responding to economic fluctuations and whether policy should be conducted by rule or by discretion. There are many arguments on both sides of these questions. Perhaps the only clear conclusion is that there is no simple and compelling case for any particular view of macroeconomic policy. In the end, you must weigh the various arguments, both economic and political, and decide for yourself what kind of role the government should play in trying to stabilize the economy.

For better or worse, economists play a key role in the formulation of economic policy. Because the economy is complex, this role is often difficult. Yet it is also inevitable. Economists cannot sit back and wait until our knowledge of the economy has been perfected before giving advice. In the meantime, someone must advise economic policymakers. That job, difficult as it sometimes is, falls to economists.

The role of economists in the policymaking process goes beyond giving advice to policymakers. Even economists cloistered in academia influence policy indirectly through their research and writing. In the conclusion of *The General Theory*, John Maynard Keynes wrote that

the ideas of economists and political philosophers, both when they are right and when they are wrong, are more powerful than is commonly understood. Indeed, the world is ruled by little else. Practical men, who believe themselves to be quite exempt from intellectual influences, are usually the slaves of some defunct economist. Madmen in authority, who hear voices in the air, are distilling their frenzy from some academic scribbler of a few years back.

This is as true today as it was when Keynes wrote it in 1936—except now that academic scribbler is often Keynes himself.

⁸ For a more complete presentation of these findings and references to the large literature on central-bank independence, see Alberto Alesina and Lawrence H. Summers, “Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence,” *Journal of Money, Credit, and Banking* 25 (May 1993): 151–162. For a study that questions the link between inflation and central-bank independence, see Marta Campillo and Jeffrey A. Miron, “Why Does Inflation Differ Across Countries?” in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy* (Chicago: University of Chicago Press, 1997), 335–362.

Summary

1. Advocates of active policy view the economy as subject to frequent shocks that will lead to unnecessary fluctuations in output and employment unless monetary or fiscal policy responds. Many believe that economic policy has been successful in stabilizing the economy.
2. Advocates of passive policy argue that because monetary and fiscal policies work with long and variable lags, attempts to stabilize the economy are likely to end up being destabilizing. In addition, they believe that our present understanding of the economy is too limited to be useful in formulating successful stabilization policy and that inept policy is a frequent source of economic fluctuations.
3. Advocates of discretionary policy argue that discretion gives more flexibility to policymakers in responding to various unforeseen situations.
4. Advocates of policy rules argue that the political process cannot be trusted. They believe that politicians make frequent mistakes in conducting economic policy and sometimes use economic policy for their own political ends. In addition, advocates of policy rules argue that a commitment to a fixed policy rule is necessary to solve the problem of time inconsistency.

KEY CONCEPTS

Inside and outside lags

Lucas critique

Time inconsistency

Automatic stabilizers

Political business cycle

Monetarists

Leading indicators

QUESTIONS FOR REVIEW

1. What are the inside lag and the outside lag? Which has the longer inside lag—monetary or fiscal policy? Which has the longer outside lag? Why?
2. Why would more accurate economic forecasting make it easier for policymakers to stabilize the economy? Describe two ways economists try to forecast developments in the economy.
3. Describe the Lucas critique.
4. How does a person's interpretation of macroeconomic history affect his view of macroeconomic policy?
5. What is meant by the "time inconsistency" of economic policy? Why might policymakers be tempted to renege on an announcement they made earlier? In this situation, what is the advantage of a policy rule?
6. List three policy rules that the Fed might follow. Which of these would you advocate? Why?

PROBLEMS AND APPLICATIONS

1. Suppose that the tradeoff between unemployment and inflation is determined by the Phillips curve:

$$u = u^n - \alpha(\pi - \pi^e),$$

where u denotes the unemployment rate, u^n the natural rate, π the rate of inflation, and π^e the expected rate of inflation. In addition, suppose that the Democratic party always follows a policy of high money growth and the Republican party always follows a policy of low money growth. What “political business cycle” pattern of inflation and unemployment would you predict under the following conditions?

- a. Every four years, one of the parties takes control based on a random flip of a coin. (*Hint:* What will expected inflation be prior to the election?)
 - b. The two parties take turns.
2. When cities pass laws limiting the rent landlords can charge on apartments, the laws usually apply to existing buildings and exempt any buildings not yet built. Advocates of rent control argue that this exemption ensures that rent control does not discourage the construction of new housing. Evaluate this argument in light of the time-inconsistency problem.
3. Go to the Web site of the Federal Reserve (www.federalreserve.gov). Find and read a press release, congressional testimony, or a report about recent monetary policy. What does it say? What is the Fed doing? Why? What do you think about the Fed’s recent policy decisions?

APPENDIX

Time Inconsistency and the Tradeoff Between Inflation and Unemployment

In this appendix, we examine more analytically the time-inconsistency argument for rules rather than discretion. This material is relegated to an appendix because we will need to use some calculus.⁹

Suppose that the Phillips curve describes the relationship between inflation and unemployment. Letting u denote the unemployment rate, u^n the natural rate of unemployment, π the rate of inflation, and π^e the expected rate of inflation, unemployment is determined by

$$u = u^n - \alpha(\pi - \pi^e).$$

Unemployment is low when inflation exceeds expected inflation and high when inflation falls below expected inflation. The parameter α determines how much unemployment responds to surprise inflation.

For simplicity, suppose also that the Fed chooses the rate of inflation. Of course, more realistically, the Fed controls inflation only imperfectly through its control of the money supply. But for purposes of illustration, it is useful to assume that the Fed can control inflation perfectly.

The Fed likes low unemployment and low inflation. Suppose that the cost of unemployment and inflation, as perceived by the Fed, can be represented as

$$L(u, \pi) = u + \gamma\pi^2,$$

where the parameter γ represents how much the Fed dislikes inflation relative to unemployment. $L(u, \pi)$ is called the *loss function*. The Fed's objective is to make the loss as small as possible.

Having specified how the economy works and the Fed's objective, let's compare monetary policy made under a fixed rule and under discretion.

First, consider policy under a fixed rule. A rule commits the Fed to a particular level of inflation. As long as private agents understand that the Fed is committed to this rule, the expected level of inflation will be the level the Fed is committed to produce. Because expected inflation equals actual inflation ($\pi^e = \pi$), unemployment will be at its natural rate ($u = u^n$).

What is the optimal rule? Because unemployment is at its natural rate regardless of the level of inflation legislated by the rule, there is no benefit to having any inflation at all. Therefore, the optimal fixed rule requires that the Fed produce zero inflation.

⁹ The material in this appendix is derived from Finn E. Kydland and Edward C. Prescott, "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85 (June 1977): 473–492; and Robert J. Barro and David Gordon, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy* 91 (August 1983): 589–610.

Second, consider discretionary monetary policy. Under discretion, the economy works as follows:

1. Private agents form their expectations of inflation π^e .
2. The Fed chooses the actual level of inflation π .
3. Based on expected and actual inflation, unemployment is determined.

Under this arrangement, the Fed minimizes its loss $L(u, \pi)$ subject to the constraint that the Phillips curve imposes. When making its decision about the rate of inflation, the Fed takes expected inflation as already determined.

To find what outcome we would obtain under discretionary policy, we must examine what level of inflation the Fed would choose. By substituting the Phillips curve into the Fed's loss function, we obtain

$$L(u, \pi) = u^n - \alpha(\pi - \pi^e) + \gamma\pi^2.$$

Notice that the Fed's loss is negatively related to unexpected inflation (the second term in the equation) and positively related to actual inflation (the third term). To find the level of inflation that minimizes this loss, differentiate with respect to π to obtain

$$dL/d\pi = -\alpha + 2\gamma\pi.$$

The loss is minimized when this derivative equals zero. Solving for π , we get

$$\pi = \alpha/(2\gamma).$$

Whatever level of inflation private agents expected, this is the "optimal" level of inflation for the Fed to choose. Of course, rational private agents understand the objective of the Fed and the constraint that the Phillips curve imposes. They therefore expect that the Fed will choose this level of inflation. Expected inflation equals actual inflation [$\pi^e = \pi = \alpha/(2\gamma)$], and unemployment equals its natural rate ($u = u^n$).

Now compare the outcome under optimal discretion to the outcome under the optimal rule. In both cases, unemployment is at its natural rate. Yet discretionary policy produces more inflation than does policy under the rule. *Thus, optimal discretion is worse than the optimal rule.* This is true even though the Fed under discretion was attempting to minimize its loss, $L(u, \pi)$.

At first it may seem strange that the Fed can achieve a better outcome by being committed to a fixed rule. Why can't the Fed with discretion mimic the Fed committed to a zero-inflation rule? The answer is that the Fed is playing a game against private decisionmakers who have rational expectations. Unless it is committed to a fixed rule of zero inflation, the Fed cannot get private agents to expect zero inflation.

Suppose, for example, that the Fed simply announces that it will follow a zero-inflation policy. Such an announcement by itself cannot be credible. After private agents have formed their expectations of inflation, the Fed has the incentive to renege on its announcement in order to decrease unemployment. (As we have

just seen, once expectations are given, the Fed's optimal policy is to set inflation at $\pi = \alpha/(2\gamma)$, regardless of π^e .) Private agents understand the incentive to renege and therefore do not believe the announcement in the first place.

This theory of monetary policy has an important corollary. Under one circumstance, the Fed with discretion achieves the same outcome as the Fed committed to a fixed rule of zero inflation. If the Fed dislikes inflation much more than it dislikes unemployment (so that γ is very large), inflation under discretion is near zero, because the Fed has little incentive to inflate. This finding provides some guidance to those who have the job of appointing central bankers. An alternative to imposing a fixed rule is to appoint an individual with a fervent distaste for inflation. Perhaps this is why even liberal politicians (Jimmy Carter, Bill Clinton) who are more concerned about unemployment than inflation sometimes appoint conservative central bankers (Paul Volcker, Alan Greenspan) who are more concerned about inflation.

MORE PROBLEMS AND APPLICATIONS

1. In the 1970s in the United States, the inflation rate and the natural rate of unemployment both rose. Let's use this model of time inconsistency to examine this phenomenon. Assume that policy is discretionary.
 - a. In the model as developed so far, what happens to the inflation rate when the natural rate of unemployment rises?
 - b. Let's now change the model slightly by supposing that the Fed's loss function is quadratic in both inflation and unemployment. That is,

$$L(u, \pi) = u^2 + \gamma\pi^2.$$

Follow steps similar to those in the text to solve for the inflation rate under discretionary policy.

- c. Now what happens to the inflation rate when the natural rate of unemployment rises?
- d. In 1979, President Jimmy Carter appointed the conservative central banker Paul Volcker to head the Federal Reserve. According to this model, what should have happened to inflation and unemployment?

Government Debt

Blessed are the young, for they shall inherit the national debt.

— Herbert Hoover

When a government spends more than it collects in taxes, it borrows from the private sector to finance the budget deficit. The accumulation of past borrowing is the government debt. Debate about the appropriate amount of government debt in the United States is as old as the country itself. Alexander Hamilton believed that “a national debt, if it is not excessive, will be to us a national blessing,” whereas James Madison argued that “a public debt is a public curse.” Indeed, the location of the nation’s capital was chosen as part of a deal in which the federal government assumed the Revolutionary War debts of the states: because the Northern states had larger outstanding debts, the capital was located in the South.

Although attention to the government debt has waxed and waned over the years, it was especially intense during the last two decades of the twentieth century. Beginning in the early 1980s, the U.S. federal government began running large budget deficits—in part because of increased spending and in part because of reduced taxes. As a result, the government debt expressed as a percentage of GDP roughly doubled from 26 percent in 1980 to 50 percent in 1995. By the late 1990s, the budget deficit had come under control and had even turned into a budget surplus. Policymakers then turned to the question of how rapidly the debt should be paid off.

The large increase in government debt from 1980 to 1995 is without precedent in U.S. history. Government debt most often rises in periods of war or depression, but the United States experienced neither during this time. Not surprisingly, the episode sparked a renewed interest among economists and policymakers in the economic effects of government debt. Some view the large budget deficits of the 1980s and 1990s as the worst mistake of economic policy since the Great Depression, whereas others think that the deficits matter very little. This chapter considers various facets of this debate.

We begin by looking at the numbers. Section 15-1 examines the size of the U.S. government debt, comparing it to the debt of other countries and to the debt that the United States has had during its own past. It also takes a brief look at what the future may hold. Section 15-2 discusses why measuring changes in

government indebtedness is not as straightforward as it might seem. Indeed, some economists argue that traditional measures are so misleading that they should be completely ignored.

We then look at how government debt affects the economy. Section 15-3 describes the traditional view of government debt, according to which government borrowing reduces national saving and crowds out capital accumulation. This view is held by most economists and has been implicit in the discussion of fiscal policy throughout this book. Section 15-4 discusses an alternative view, called *Ricardian equivalence*, which is held by a small but influential minority of economists. According to the Ricardian view, government debt does not influence national saving and capital accumulation. As we will see, the debate between the traditional and Ricardian views of government debt arises from disagreements over how consumers respond to the government's debt policy.

Section 15-5 then looks at other facets of the debate over government debt. It begins by discussing whether the government should try to always balance its budget and, if not, when a budget deficit or surplus is desirable. It also examines the effects of government debt on monetary policy, the political process, and the role of a country in the world economy.

15-1 The Size of the Government Debt

Let's begin by putting the government debt in perspective. In 2001, the debt of the U.S. federal government was \$3.2 trillion. If we divide this number by 276 million, the number of people in the United States, we find that each person's share of the government debt was about \$11,600. Obviously, this is not a trivial number—few people sneeze at \$11,600. Yet if we compare this debt to

table 15-1

How Indebted Are the World's Governments?

Country	Government Debt as a Percentage of GDP	Country	Government Debt as a Percentage of GDP
Japan	119	Ireland	54
Italy	108	Spain	53
Belgium	105	Finland	51
Canada	101	Sweden	49
Greece	100	Germany	46
Denmark	67	Austria	40
United Kingdom	64	Netherlands	27
United States	62	Australia	26
France	58	Norway	24
Portugal	55		

Source: OECD Economic Outlook. Figures are based on estimates of gross government debt and GDP for 2001.


the roughly \$1 million a typical person will earn over his or her working life, the government debt does not look like the catastrophe it is sometimes made out to be.

One way to judge the size of a government's debt is to compare it to the amount of debt other countries have accumulated. Table 15-1 shows the amount of government debt for 19 major countries expressed as a percentage of each country's GDP. On the top of the list are the heavily indebted countries of Japan and Italy, which have accumulated a debt that exceeds annual GDP. At the bottom are Norway and Australia, which have accumulated relatively small debts. The United States is in the middle of the pack. By international standards, the U.S. government is neither especially profligate nor especially frugal.

Over the course of U.S. history, the indebtedness of the federal government has varied substantially. Figure 15-1 shows the ratio of the federal debt to GDP since 1791. The government debt, relative to the size of the economy, varies from close to zero in the 1830s to a maximum of 107 percent of GDP in 1945.

Historically, the primary cause of increases in the government debt is war. The debt-GDP ratio rises sharply during major wars and falls slowly during peacetime. Many economists think that this historical pattern is the appropriate

figure 15-1



The Ratio of Government Debt to GDP Since 1790 The U.S. federal government debt held by the public, relative to the size of the U.S. economy, rises sharply during wars and declines slowly during peacetime. The exception is the period since 1980, when the debt-GDP ratio rose without the occurrence of a major military conflict.

Source: U.S. Department of Treasury, U.S. Department of Commerce, and T.S. Berry, "Production and Population Since 1789," Bostwick Paper No. 6, Richmond, 1988.

way to run fiscal policy. As we discuss more fully later in this chapter, deficit financing of wars appears optimal for reasons of both tax smoothing and generational equity. One instance of a large increase in government debt in peacetime occurred during the 1980s and early 1990s, when the federal government ran substantial budget deficits. Many economists have criticized this increase in government debt as imposing a burden on future generations without justification.

During the middle of the 1990s, the U.S. federal government started to get its budget deficit under control. A combination of tax hikes, spending cuts, and rapid economic growth caused the ratio of debt to GDP to stabilize and then decline. Recent experience has tempted some observers to think that exploding government debt is a thing of the past. But as the next case study suggests, the worst may be yet to come.

CASE STUDY

The Fiscal Future: Good News and Bad News

What does the future hold for fiscal policymakers? Economic forecasting is far from precise, and it is easy to be cynical about economic predictions. But good policy cannot be made if policymakers only look backwards. As a result, economists in the Congressional Budget Office (CBO) and other government agencies are always trying to look ahead to see what problems and opportunities are likely to develop.

When George W. Bush moved into the White House in 2001, the fiscal picture facing the U.S. government was mixed. In particular, it depended on how far one looked ahead.

Over a ten- or twenty-year horizon, the picture looked good. The U.S. federal government was running a large budget surplus. As a percentage of GDP, the projected surplus for 2001 was the largest since 1948. Moreover, the surplus was expected to grow even larger over time. The surplus was large enough so that, without any policy changes, the government debt would be paid off by 2008.

These surpluses arose from various sources. The elder George Bush had signed a tax increase in 1990, and Bill Clinton had signed another in 1993. Because of these tax hikes, federal tax revenue as a percentage of GDP reached its highest level since World War II. Then, in the late 1990s, productivity accelerated, most likely because of advances in information technology. The high growth in incomes led to rising tax revenue, which pushed the federal government's budget from deficit to surplus.

A debate arose over how to respond to the budget surplus. The government could use the large projected surpluses to repay debt, increase spending, cut taxes, or some combination of these. The new Republican president George W. Bush advocated a tax cut of \$1.6 trillion over 10 years, which was about one-fourth of the projected surpluses. Democrats in Congress argued for a smaller tax cut and greater government spending. The end result was a compromise bill that cut taxes by a bit less than Bush had advocated.

While the 10-year horizon looked rosy, the longer-term fiscal picture was more troublesome. The problem was demographic. Advances in medical technology have been increasing life expectancy, while improvements in birth-control techniques and changing social norms have reduced the number of children people have. Because of these developments, the elderly are becoming a larger share of the population. In 1990, there were 21 elderly for every 100 people of working age (ages 20 to 64); this figure is projected to rise to 36 by the year 2030. Such a demographic change has profound implications for fiscal policy. About one-third of the budget of the U.S. federal government is devoted to pensions and health care for the elderly. As more people become eligible for these “entitlements,” as they are sometimes called, government spending automatically rises over time, pushing the budget toward deficit.

The magnitude of these budgetary pressures was documented in a CBO report released in October 2000. According to the CBO, if no changes in fiscal policy are enacted, the government debt as a percentage of GDP will start rising around 2030 and reach historic highs around 2060. At that point, the government’s budget will spiral out of control.¹

Of course, all economic forecasts need to be greeted with a bit of skepticism, especially those that try to look ahead half a century. Shocks to the economy can alter the government’s revenue and spending. In fact, only months after moving into the White House, George W. Bush saw the fiscal picture start to change. First, the economic slowdown in 2001 reduced tax revenue. Then, the terrorist attacks in September 2001 induced an increase in government spending. Both developments reduced the projected near-term government surpluses. As this book was going to press, there was great uncertainty about future government spending and the rate of technological advance—two key determinants of the fiscal situation.

Yet one thing is clear: the elderly are making up a larger share of the population, and this fact will shape the fiscal challenges in the decades ahead.

15-2 Problems in Measurement

The government budget deficit equals government spending minus government revenue, which in turn equals the amount of new debt the government needs to issue to finance its operations. This definition may sound simple enough, but in fact debates over fiscal policy sometimes arise over how the budget deficit should be measured. Some economists believe that the deficit as currently measured is not a good indicator of the stance of fiscal policy. That is, they believe that the budget deficit does not accurately gauge either the impact of fiscal policy on today’s economy or the burden being placed on future generations of taxpayers. In this section we discuss four problems with the usual measure of the budget deficit.

¹ Congressional Budget Office, *The Long-Term Budget Outlook*, October 2000.

Measurement Problem 1: Inflation

The least controversial of the measurement issues is the correction for inflation. Almost all economists agree that the government's indebtedness should be measured in real terms, not in nominal terms. The measured deficit should equal the change in the government's real debt, not the change in its nominal debt.

The budget deficit as commonly measured, however, does not correct for inflation. To see how large an error this induces, consider the following example. Suppose that the real government debt is not changing; in other words, in real terms, the budget is balanced. In this case, the nominal debt must be rising at the rate of inflation. That is,

$$\Delta D/D = \pi,$$

where π is the inflation rate and D is the stock of government debt. This implies

$$\Delta D = \pi D.$$

The government would look at the change in the nominal debt ΔD and would report a budget deficit of πD . Hence, most economists believe that the reported budget deficit is overstated by the amount πD .

We can make the same argument in another way. The deficit is government expenditure minus government revenue. Part of expenditure is the interest paid on the government debt. Expenditure should include only the real interest paid on the debt rD , not the nominal interest paid iD . Because the difference between the nominal interest rate i and the real interest rate r is the inflation rate π , the budget deficit is overstated by πD .

This correction for inflation can be large, especially when inflation is high, and it can often change our evaluation of fiscal policy. For example, in 1979, the federal government reported a budget deficit of \$28 billion. Inflation was 8.6 percent, and the government debt held at the beginning of the year by the public (excluding the Federal Reserve) was \$495 billion. The deficit was therefore overstated by

$$\begin{aligned}\pi D &= 0.086 \times \$495 \text{ billion} \\ &= \$43 \text{ billion.}\end{aligned}$$

Corrected for inflation, the reported budget deficit of \$28 billion turns into a budget surplus of \$15 billion! In other words, even though nominal government debt was rising, real government debt was falling.

Measurement Problem 2: Capital Assets

Many economists believe that an accurate assessment of the government's budget deficit requires accounting for the government's assets as well as its liabilities. In particular, when measuring the government's overall indebtedness, we should subtract government assets from government debt. Therefore, the budget deficit should be measured as the change in debt minus the change in assets.

Certainly, individuals and firms treat assets and liabilities symmetrically. When a person borrows to buy a house, we do not say that he is running a budget deficit. Instead, we offset the increase in assets (the house) against the increase in debt (the mortgage) and record no change in net wealth. Perhaps we should treat the government's finances the same way.

A budget procedure that accounts for assets as well as liabilities is called **capital budgeting**, because it takes into account changes in capital. For example, suppose that the government sells one of its office buildings or some of its land and uses the proceeds to reduce the government debt. Under current budget procedures, the reported deficit would be lower. Under capital budgeting, the revenue received from the sale would not lower the deficit, because the reduction in debt would be offset by a reduction in assets. Similarly, under capital budgeting, government borrowing to finance the purchase of a capital good would not raise the deficit.

The major difficulty with capital budgeting is that it is hard to decide which government expenditures should count as capital expenditures. For example, should the interstate highway system be counted as an asset of the government? If so, what is its value? What about the stockpile of nuclear weapons? Should spending on education be treated as expenditure on human capital? These difficult questions must be answered if the government is to adopt a capital budget.

Economists and policymakers disagree about whether the federal government should use capital budgeting. (Many state governments already use it.) Opponents of capital budgeting argue that, although the system is superior in principle to the current system, it is too difficult to implement in practice. Proponents of capital budgeting argue that even an imperfect treatment of capital assets would be better than ignoring them altogether.

Measurement Problem 3: Uncounted Liabilities

Some economists argue that the measured budget deficit is misleading because it excludes some important government liabilities. For example, consider the pensions of government workers. These workers provide labor services to the government today, but part of their compensation is deferred to the future. In essence, these workers are providing a loan to the government. Their future pension benefits represent a government liability not very different from government debt. Yet this liability is not included as part of the government debt, and the accumulation of this liability is not included as part of the budget deficit. According to some estimates, this implicit liability is almost as large as the official government debt.

Similarly, consider the Social Security system. In some ways, the system is like a pension plan. People pay some of their income into the system when young and expect to receive benefits when old. Perhaps accumulated future Social Security benefits should be included in the government's liabilities. Estimates suggest that the government's future Social Security liabilities (less future Social Security taxes) are more than three times the government debt as officially measured.

One might argue that Social Security liabilities are different from government debt because the government can change the laws determining Social Security

benefits. Yet, in principle, the government could always choose not to repay all of its debt: the government honors its debt only because it chooses to do so. Promises to pay the holders of government debt may not be fundamentally different from promises to pay the future recipients of Social Security.

A particularly difficult form of government liability to measure is the *contingent liability*—the liability that is due only if a specified event occurs. For example, the government guarantees many forms of private credit, such as student loans, mortgages for low- and moderate-income families, and deposits in banks and savings-and-loan institutions. If the borrower repays the loan, the government pays nothing; if the borrower defaults, the government makes the repayment. When the government provides this guarantee, it undertakes a liability contingent on the borrower's default. Yet this contingent liability is not reflected in the budget deficit, in part because it is not clear what dollar value to attach to it.

Measurement Problem 4: The Business Cycle

Many changes in the government's budget deficit occur automatically in response to a fluctuating economy. For example, when the economy goes into a recession, incomes fall, so people pay less in personal income taxes. Profits fall, so corporations pay less in corporate income taxes. More people become eligible for government assistance, such as welfare and unemployment insurance, so government spending rises. Even without any change in the laws governing taxation and spending, the budget deficit increases.

These automatic changes in the deficit are not errors in measurement, because the government truly borrows more when a recession depresses tax revenue and boosts government spending. But these changes do make it more difficult to use the deficit to monitor changes in fiscal policy. That is, the deficit can rise or fall either because the government has changed policy or because the economy has changed direction. For some purposes, it would be good to know which is occurring.

To solve this problem, the government calculates a **cyclically adjusted budget deficit** (sometimes called the *full-employment budget deficit*). The cyclically adjusted deficit is based on estimates of what government spending and tax revenue would be if the economy were operating at its natural rate of output and employment. The cyclically adjusted deficit is a useful measure because it reflects policy changes but not the current stage of the business cycle.

Summing Up

Economists differ in the importance they place on these measurement problems. Some believe that the problems are so severe that the measured budget deficit is almost meaningless. Most take these measurement problems seriously but still view the measured budget deficit as a useful indicator of fiscal policy.

The undisputed lesson is that to evaluate fully what fiscal policy is doing, economists and policymakers must look at more than only the measured budget deficit. And, in fact, they do. The budget documents prepared annually by the Office of Management and Budget contain much detailed information about the government's finances, including data on capital expenditures and credit programs.

No economic statistic is perfect. Whenever we see a number reported in the media, we need to know what it is measuring and what it is leaving out. This is especially true for data on government debt and budget deficits.

CASE STUDY

Generational Accounting

One harsh critic of current measures of the budget deficit is economist Laurence Kotlikoff. Kotlikoff argues that the budget deficit is like the fabled emperor who wore no clothes: everyone should plainly see the problem, but no one is willing to admit to it. He writes, “On the conceptual level, the budget deficit is intellectually bankrupt. On the practical level, there are so many official deficits that ‘balanced budget’ has lost any true meaning.” He sees an “urgent need to switch from an outdated, misleading, and fundamentally noneconomic measure of fiscal policy, namely the budget deficit, to generational accounting.”

Generational accounting, Kotlikoff’s new way to gauge the influence of fiscal policy, is based on the idea that a person’s economic well-being depends on his or her lifetime income. (This idea is founded on Modigliani’s life-cycle theory of consumer behavior, which we examine in Chapter 16.) When evaluating fiscal policy, therefore, we should not be concerned with taxes or spending in any single year. Instead, we should look at the taxes paid, and transfers received, by people over their entire lives. Generational accounts measure the impact of fiscal policy on the lifetime incomes of different generations.

Generational accounts tell a very different story than the budget deficit about the history of U.S. fiscal policy. In the early 1980s, the U.S. government cut taxes, beginning a long period of large budget deficits. Most commentators claim that older generations benefited at the expense of younger generations during this period, because the young inherited the government debt. Kotlikoff agrees that these tax cuts raised the burden on the young, but he claims that this standard analysis ignores the impact of many other policy changes. His generational accounts show that the young were hit even harder during the 1950s, 1960s, and 1970s. During these years, the government raised Social Security benefits for the elderly and financed the higher spending by taxing the working-age population. This policy redistributed income away from the young, even though it did not affect the budget deficit. During the 1980s, Social Security reforms reversed this trend, benefiting younger generations.

Despite Kotlikoff’s advocacy, generational accounting is not likely to replace the budget deficit. This alternative system also has flaws. For example, to calculate the total tax burden on different generations, one needs to make assumptions about future policy, which are open to dispute. Nonetheless, generational accounting offers a useful perspective in the debate over fiscal policy.²

² Laurence J. Kotlikoff, *Generational Accounting: Knowing Who Pays, and When, for What We Spend* (New York: The Free Press, 1992). For an appraisal of the book, see David M. Cutler, Book Review, *National Tax Journal* 56 (March 1993): 61–67. See also the symposium on generational accounting in the Winter 1994 issue of the *Journal of Economic Perspectives*.

15-3 The Traditional View of Government Debt

Imagine that you are an economist working for the Congressional Budget Office (CBO). You receive a letter from the chair of the Senate Budget Committee:

Dear CBO Economist:

Congress is about to consider the president's request to cut all taxes by 20 percent. Before deciding whether to endorse the request, my committee would like your analysis. We see little hope of reducing government spending, so the tax cut would mean an increase in the budget deficit. How would the tax cut and budget deficit affect the economy and the economic well-being of the country?

Sincerely,
Committee Chair

Before responding to the senator, you open your favorite economics textbook—to this one, of course—to see what the models predict for such a change in fiscal policy.

To analyze the long-run effects of this policy change, you turn to the models in Chapters 3 through 8. The model in Chapter 3 shows that a tax cut stimulates consumer spending and reduces national saving. The reduction in saving raises the interest rate, which crowds out investment. The Solow growth model introduced in Chapter 7 shows that lower investment eventually leads to a lower steady-state capital stock and a lower level of output. Because we concluded in Chapter 8 that the U.S. economy has less capital than in the Golden Rule steady state (the steady state with maximum consumption), the fall in steady-state capital means lower consumption and reduced economic well-being.

To analyze the short-run effects of the policy change, you turn to the *IS-LM* model in Chapters 10 and 11. This model shows that a tax cut stimulates consumer spending, which implies an expansionary shift in the *IS* curve. If there is no change in monetary policy, the shift in the *IS* curve leads to an expansionary shift in the aggregate demand curve. In the short run, when prices are sticky, the expansion in aggregate demand leads to higher output and lower unemployment. Over time, as prices adjust, the economy returns to the natural rate of output, and the higher aggregate demand results in a higher price level.

To see how international trade affects your analysis, you turn to the open-economy models in Chapters 5 and 12. The model in Chapter 5 shows that when national saving falls, people start financing investment by borrowing from abroad, causing a trade deficit. Although the inflow of capital from abroad lessens the effect of the fiscal-policy change on U.S. capital accumulation, the United States becomes indebted to foreign countries. The fiscal-policy change also causes the dollar to appreciate, which makes foreign goods cheaper in the United States and domestic goods more expensive abroad. The Mundell-Fleming model in Chapter 12 shows that the appreciation of the dollar and the resulting fall in net exports reduce the short-run expansionary impact of the fiscal change on output and employment.

With all these models in mind, you draft a response:

Dear Senator:

A tax cut financed by government borrowing would have many effects on the economy. The immediate impact of the tax cut would be to stimulate consumer spending. Higher consumer spending affects the economy in both the short run and the long run.

In the short run, higher consumer spending would raise the demand for goods and services and thus raise output and employment. Interest rates would also rise, however, as investors competed for a smaller flow of saving. Higher interest rates would discourage investment and would encourage capital to flow in from abroad. The dollar would rise in value against foreign currencies, and U.S. firms would become less competitive in world markets.

In the long run, the smaller national saving caused by the tax cut would mean a smaller capital stock and a greater foreign debt. Therefore, the output of the nation would be smaller, and a greater share of that output would be owed to foreigners.

The overall effect of the tax cut on economic well-being is hard to judge. Current generations would benefit from higher consumption and higher employment, although inflation would likely be higher as well. Future generations would bear much of the burden of today's budget deficits: they would be born into a nation with a smaller capital stock and a larger foreign debt.

Your faithful servant,
CBO Economist

The senator replies:

Dear CBO Economist:

Thank you for your letter. It made sense to me. But yesterday my committee heard testimony from a prominent economist who called herself a “Ricardian” and who reached quite a different conclusion. She said that a tax cut by itself would not stimulate consumer spending. She concluded that the budget deficit would therefore not have all the effects you listed. What’s going on here?

Sincerely,
Committee Chair

After studying the next section, you write back to the senator, explaining in detail the debate over Ricardian equivalence.

15-4 The Ricardian View of Government Debt

The traditional view of government debt presumes that when the government cuts taxes and runs a budget deficit, consumers respond to their higher after-tax income by spending more. An alternative view, called **Ricardian equivalence**, questions this presumption. According to the Ricardian view, consumers are forward-looking and, therefore, base their spending not only on their current income but also on their expected future income. As we explore more fully in Chapter 16, the forward-looking consumer is at the heart of many modern theories of consumption. The Ricardian view of government debt applies the logic of the forward-looking consumer to analyze the effects of fiscal policy.

The Basic Logic of Ricardian Equivalence

Consider the response of a forward-looking consumer to the tax cut that the Senate Budget Committee is considering. The consumer might reason as follows:

The government is cutting taxes without any plans to reduce government spending. Does this policy alter my set of opportunities? Am I richer because of this tax cut? Should I consume more?

Maybe not. The government is financing the tax cut by running a budget deficit. At some point in the future, the government will have to raise taxes to pay off the debt and accumulated interest. So the policy really represents a tax cut today coupled with a tax hike in the future. The tax cut merely gives me transitory income that eventually will be taken back. I am not any better off, so I will leave my consumption unchanged.

The forward-looking consumer understands that government borrowing today means higher taxes in the future. A tax cut financed by government debt does not reduce the tax burden; it merely reschedules it. It therefore should not encourage the consumer to spend more.

One can view this argument another way. Suppose that the government borrows \$1,000 from the typical citizen to give that citizen a \$1,000 tax cut. In essence, this policy is the same as giving the citizen a \$1,000 government bond as a gift. One side of the bond says, “The government owes you, the bondholder, \$1,000 plus interest.” The other side says, “You, the taxpayer, owe the government \$1,000 plus interest.” Overall, the gift of a bond from the government to the typical citizen does not make the citizen richer or poorer, because the value of the bond is offset by the value of the future tax liability.

The general principle is that government debt is equivalent to future taxes, and if consumers are sufficiently forward-looking, future taxes are equivalent to current taxes. Hence, financing the government by debt is equivalent to financing it by taxes. This view is called *Ricardian equivalence* after the famous nineteenth-century economist David Ricardo, because he first noted the theoretical argument.

The implication of Ricardian equivalence is that a debt-financed tax cut leaves consumption unaffected. Households save the extra disposable income to pay the future tax liability that the tax cut implies. This increase in private saving exactly offsets the decrease in public saving. National saving—the sum of private and public saving—remains the same. The tax cut therefore has none of the effects that the traditional analysis predicts.

The logic of Ricardian equivalence does not mean that all changes in fiscal policy are irrelevant. Changes in fiscal policy do influence consumer spending if they influence present or future government purchases. For example, suppose that the government cuts taxes today because it plans to reduce government purchases in the future. If the consumer understands that this tax cut does not require an increase in future taxes, he feels richer and raises his consumption. But note that it is the reduction in government purchases, rather than the reduction in taxes, that stimulates consumption: the announcement

of a future reduction in government purchases would raise consumption today even if current taxes were unchanged, because it would imply lower taxes at some time in the future.

Consumers and Future Taxes

The essence of the Ricardian view is that when people choose their consumption, they rationally look ahead to the future taxes implied by government debt. But how forward-looking are consumers? Defenders of the traditional view of government debt believe that the prospect of future taxes does not have as large an influence on current consumption as the Ricardian view assumes. Here are some of their arguments.³

Myopia Proponents of the Ricardian view of fiscal policy assume that people are rational when making decisions such as choosing how much of their income to consume and how much to save. When the government borrows to pay for current spending, rational consumers look ahead to the future taxes required to support this debt. Thus, the Ricardian view presumes that people have substantial knowledge and foresight.

One possible argument for the traditional view of tax cuts is that people are shortsighted, perhaps because they do not fully comprehend the implications of government budget deficits. It is possible that some people follow simple and not fully rational rules of thumb when choosing how much to save. Suppose, for example, that a person acts on the assumption that future taxes will be the same as current taxes. This person will fail to take account of future changes in taxes required by current government policies. A debt-financed tax cut will lead this person to believe that his lifetime income has increased, even if it hasn't. The tax cut will therefore lead to higher consumption and lower national saving.

Borrowing Constraints The Ricardian view of government debt assumes that consumers base their spending not only on current income but on their lifetime income, which includes both current and expected future income. According to the Ricardian view, a debt-financed tax cut increases current income, but it does not alter lifetime income or consumption. Advocates of the traditional view of government debt argue that current income is more important than lifetime income for those consumers who face binding borrowing constraints. A *borrowing constraint* is a limit on how much an individual can borrow from banks or other financial institutions.

A person who would like to consume more than his current income—perhaps because he expects higher income in the future—has to do so by borrowing. If he cannot borrow to finance current consumption, or can borrow only a

³ For a survey of the debate over Ricardian equivalence, see Douglas Bernheim, “Ricardian Equivalence: An Evaluation of Theory and Evidence,” *NBER Macroeconomics Annual* (1987): 263–303. See also the symposium on budget deficits in the Spring 1989 issue of the *Journal of Economic Perspectives*.

limited amount, his current income determines his spending, regardless of what his lifetime income might be. In this case, a debt-financed tax cut raises current income and thus consumption, even though future income is lower. In essence, when the government cuts current taxes and raises future taxes, it is giving taxpayers a loan. For a person who wanted to obtain a loan but was unable to, the tax cut expands his opportunities and stimulates consumption.

CASE STUDY

George Bush's Withholding Experiment

In early 1992, President George Bush pursued a novel policy to deal with the lingering recession in the United States. By executive order, he lowered the amount of income taxes that were being withheld from workers' paychecks. The order did not reduce the amount of taxes that workers owed; it merely delayed payment. The higher take-home pay that workers received during 1992 was to be offset by higher tax payments, or smaller tax refunds, when income taxes were due in April 1993.

What effect would you predict for this policy? According to the logic of Ricardian equivalence, consumers should realize that their lifetime resources were unchanged and, therefore, save the extra take-home pay to meet the upcoming tax liability. Yet George Bush claimed his policy would provide "money people can use to help pay for clothing, college, or to get a new car." That is, he believed that consumers would spend the extra income, thereby stimulating aggregate demand and helping the economy recover from the recession. Bush seemed to be assuming that consumers were shortsighted or faced binding borrowing constraints.

Gauging the actual effects of this policy is difficult with aggregate data, because many other things were happening at the same time. Yet some evidence comes from a survey two economists conducted shortly after the policy was announced. The survey asked people what they would do with the extra income. Fifty-seven percent of the respondents said they would save it, use it to repay debts, or adjust their withholding in order to reverse the effect of Bush's executive order. Forty-three percent said they would spend the extra income. Thus, for this policy change, a majority of the population was planning to act as Ricardian theory posits. Nonetheless, Bush was partly right: many people planned to spend the extra income, even though they understood that the following year's tax bill would be higher.⁴

Future Generations Besides myopia and borrowing constraints, a third argument for the traditional view of government debt is that consumers expect the implied future taxes to fall not on them but on future generations. Suppose,

⁴ Matthew D. Shapiro and Joel Slemrod, "Consumer Response to the Timing of Income: Evidence From a Change in Tax Withholding," *American Economic Review* 85 (March 1995): 274–283.

for example, that the government cuts taxes today, issues 30-year bonds to finance the budget deficit, and then raises taxes in 30 years to repay the loan. In this case, the government debt represents a transfer of wealth from the next generation of taxpayers (which faces the tax hike) to the current generation of taxpayers (which gets the tax cut). This transfer raises the lifetime resources of the current generation, so it raises their consumption. In essence, a debt-financed tax cut stimulates consumption because it gives the current generation the opportunity to consume at the expense of the next generation.

Economist Robert Barro has provided a clever rejoinder to this argument to support the Ricardian view. Barro argues that because future generations are the children and grandchildren of the current generation, we should not view them as independent economic actors. Instead, he argues, the appropriate assumption is that current generations care about future generations. This altruism between generations is evidenced by the gifts that many people give their children, often in the form of bequests at the time of their deaths. The existence of bequests suggests that many people are not eager to take advantage of the opportunity to consume at their children's expense.

According to Barro's analysis, the relevant decisionmaking unit is not the individual, whose life is finite, but the family, which continues forever. In other words, an individual decides how much to consume based not only on his own income but also on the income of future members of his family. A debt-financed tax cut may raise the income an individual receives in his lifetime, but it does not raise his family's overall resources. Instead of consuming the extra income from the tax cut, the individual saves it and leaves it as a bequest to his children, who will bear the future tax liability.

We can see now that the debate over government debt is really a debate over consumer behavior. The Ricardian view assumes that consumers have a long time horizon. Barro's analysis of the family implies that the consumer's time horizon, like the government's, is effectively infinite. Yet it is possible that consumers do not look ahead to the tax liabilities of future generations. Perhaps they expect their children to be richer than they are and therefore welcome the opportunity to consume at their children's expense. The fact that many people leave zero or minimal bequests to their children is consistent with this hypothesis. For these zero-bequest families, a debt-financed tax cut alters consumption by redistributing wealth among generations.⁵



Drawing by Dave Carpenter. From the *Wall Street Journal*. Permission, Cartoon Features Syndicate.

"What's this I hear about you adults mortgaging my future?"

⁵ Robert J. Barro, "Are Government Bonds Net Wealth?" *Journal of Political Economy* 81 (1974): 1095–1117.

CASE STUDY

Why Do Parents Leave Bequests?

The debate over Ricardian equivalence is partly a debate over how different generations are linked to one another. Robert Barro's defense of the Ricardian view is based on the assumption that parents leave their children bequests because they care about them. But is altruism really the reason that parents leave bequests?

One group of economists has suggested that parents use bequests to control their children. Parents often want their children to do certain things for them, such as phoning home regularly and visiting on holidays. Perhaps parents use the implicit threat of disinheritance to induce their children to be more attentive.

To test this "strategic bequest motive," these economists examined data on how often children visit their parents. They found that the more wealthy the parent, the more often the children visit. Even more striking was another result: only wealth that can be left as a bequest induces more frequent visits. Wealth that cannot be bequeathed, such as pension wealth which reverts to the pension company in the event of an early death, does not encourage children to visit. These findings suggest that there may be more to the relationships among generations than mere altruism.⁶

Making a Choice

Having seen the traditional and Ricardian views of government debt, you should ask yourself two sets of questions.

First, with which view do you agree? If the government cuts taxes today, runs a budget deficit, and raises taxes in the future, how will the policy affect the economy? Will it stimulate consumption, as the traditional view holds? Or will consumers understand that their lifetime income is unchanged and, therefore, offset the budget deficit with higher private saving?

Second, why do you hold the view that you do? If you agree with the traditional view of government debt, what is the reason? Do consumers fail to understand that higher government borrowing today means higher taxes tomorrow? Or do they ignore future taxes, either because they are borrowing-constrained or because future taxes fall on future generations with which they do not feel an economic link? If you hold the Ricardian view, do you believe that consumers have the foresight to see that government borrowing today will result in future taxes levied on them or their descendants? Do you believe that consumers will save the extra income to offset that future tax liability?

We might hope that the evidence could help us decide between these two views of government debt. Yet when economists examine historical episodes of large budget deficits, the evidence is inconclusive. History can be interpreted in different ways.

⁶ B. Douglas Bernheim, Andrei Shleifer, and Lawrence H. Summers, "The Strategic Bequest Motive," *Journal of Political Economy* 93 (1985): 1045–1076.



Ricardo on Ricardian Equivalence

David Ricardo was a millionaire stockbroker and one of the great economists of all time. His most important contribution to the field was his 1817 book *Principles of Political Economy and Taxation*, in which he developed the theory of comparative advantage, which economists still use to explain the gains from international trade. Ricardo was also a member of the British Parliament, where he put his own theories to work and opposed the corn laws, which restricted international trade in grain.

Ricardo was interested in the alternative ways in which a government might pay for its expenditure. In an 1820 article called *Essay on the Funding System*, he considered an example of a war that cost 20 million pounds. He noted that if the interest rate were 5 percent, this expense could be financed with a one-time tax of 20 million pounds, a perpetual tax of 1 million pounds, or a tax of 1.2 million pounds for 45 years. He wrote

In point of economy, there is no real difference in either of the modes; for twenty million in one payment, one million per annum for ever, or 1,200,000 pounds for 45 years, are precisely of the same value.

Ricardo was aware that the issue involved the linkages among generations:

It would be difficult to convince a man possessed of 20,000 pounds, or any other sum, that a perpetual

payment of 50 pounds per annum was equally burdensome with a single tax of 1000 pounds. He would have some vague notion that the 50 pounds per annum would be paid by posterity, and would not be paid by him; but if he leaves his fortune to his son, and leaves it charged with this perpetual tax, where is the difference whether he leaves him 20,000 pounds with the tax, or 19,000 pounds without it?

Although Ricardo viewed these alternative methods of government finance as equivalent, he did not think other people would view them as such:

The people who pay taxes . . . do not manage their private affairs accordingly. We are apt to think that the war is burdensome only in proportion to what we are at the moment called to pay for it in taxes, without reflecting on the probable duration of such taxes.

Thus, Ricardo doubted that people were rational and farsighted enough to look ahead fully to their future tax liabilities.

As a policymaker, Ricardo took seriously the government debt. Before the British Parliament, he once declared,

This would be the happiest country in the world, and its progress in prosperity would go beyond the powers of imagination to conceive, if we got rid of two great evils—the national debt and the corn laws.

It is one of the great ironies in the history of economic thought that Ricardo rejected the theory that now bears his name!

Consider, for example, the experience of the 1980s. The large budget deficits, caused partly by the Reagan tax cut of 1981, seem to offer a natural experiment to test the two views of government debt. At first glance, this episode appears decisively to support the traditional view. The large budget deficits coincided with low national saving, high real interest rates, and a large trade deficit. Indeed, advocates of the traditional view of government debt often claim that the experience of the 1980s confirms their position.

Yet those who hold the Ricardian view of government debt interpret these events differently. Perhaps saving was low in the 1980s because people were optimistic about future economic growth—an optimism that was also reflected in a booming stock market. Or perhaps saving was low because people expected that the tax cut would eventually lead not to higher taxes but, as Reagan promised, to lower government spending. Because it is hard to rule out any of these interpretations, both views of government debt survive.

15-5 Other Perspectives on Government Debt

The policy debates over government debt have many facets. So far we have considered the traditional and Ricardian views of government debt. According to the traditional view, a government budget deficit expands aggregate demand and stimulates output in the short run but crowds out capital and depresses economic growth in the long run. According to the Ricardian view, a government budget deficit has none of these effects, because consumers understand that a budget deficit represents merely the postponement of a tax burden. With these two theories as background, we now consider several other perspectives on government debt.

Balanced Budgets Versus Optimal Fiscal Policy

In the United States, many state constitutions require the state government to run a balanced budget. A recurring topic of political debate is whether the federal Constitution should require a balanced budget for the federal government as well. Most economists oppose a strict rule requiring the government to balance its budget. There are three reasons why optimal fiscal policy may at times call for a budget deficit or surplus.

Stabilization A budget deficit or surplus can help stabilize the economy. In essence, a balanced-budget rule would revoke the automatic stabilizing powers of the system of taxes and transfers. When the economy goes into a recession, taxes automatically fall, and transfers automatically rise. Although these automatic responses help stabilize the economy, they push the budget into deficit. A strict balanced-budget rule would require that the government raise taxes or reduce spending in a recession, but these actions would further depress aggregate demand.

Tax Smoothing A budget deficit or surplus can be used to reduce the distortion of incentives caused by the tax system. As you probably learned in microeconomics courses, high tax rates impose a cost on society by discouraging economic activity. A tax on labor earnings, for instance, reduces the incentive that people have to work long hours. Because this disincentive becomes particularly large at very high tax rates, the total social cost of taxes is minimized by keeping tax rates relatively stable rather than making them high in some years and low in others. Economists call this policy *tax smoothing*. To keep tax rates smooth, a deficit is necessary in years of unusually low income (recessions) or unusually high expenditure (wars).

Intergenerational Redistribution A budget deficit can be used to shift a tax burden from current to future generations. For example, some economists argue that if the current generation fights a war to maintain freedom, future generations benefit as well and should bear some of the burden. To pass on some of the war's costs, the current generation can finance the war with a

budget deficit. The government can later retire the debt by levying taxes on the next generation.

These considerations lead most economists to reject a strict balanced-budget rule. At the very least, a rule for fiscal policy needs to take account of the recurring episodes, such as recessions and wars, during which a budget deficit is a reasonable policy response.

Fiscal Effects on Monetary Policy

In 1985, Paul Volcker told Congress that “the actual and prospective size of the budget deficit . . . heightens skepticism about our ability to control the money supply and contain inflation.” A decade later, Alan Greenspan claimed that “a substantial reduction in the long-term prospective deficit of the United States will significantly lower very long-term inflation expectations.” Both of these Fed chairmen apparently saw a link between fiscal policy and monetary policy.

We first discussed such a possibility in Chapter 4. As we saw, one way for a government to finance a budget deficit is simply to print money—a policy that leads to higher inflation. Indeed, when countries experience hyperinflation, the typical reason is that fiscal policymakers are relying on the inflation tax to pay for some of their spending. The ends of hyperinflations almost always coincide with fiscal reforms that include large cuts in government spending and therefore a reduced need for seigniorage.

In addition to this link between the budget deficit and inflation, some economists have suggested that a high level of debt might also encourage the government to create inflation. Because most government debt is specified in nominal terms, the real value of the debt falls when the price level rises. This is the usual redistribution between creditors and debtors caused by unexpected inflation—here the debtor is the government and the creditor is the private sector. But this debtor, unlike others, has access to the monetary printing press. A high level of debt might encourage the government to print money, thereby raising the price level and reducing the real value of its debts.

Despite these concerns about a possible link between government debt and monetary policy, there is little evidence that this link is important in most developed countries. In the United States, for instance, inflation was high in the 1970s, even though government debt was low relative to GDP. Monetary policymakers got inflation under control in the early 1980s, just as fiscal policymakers started running large budget deficits and increasing the government debt. Thus, although monetary policy might be driven by fiscal policy in some situations, such as during the classic hyperinflations, this situation appears not to be the norm in most countries today. There are several reasons for this. First, most governments can finance deficits by selling debt and don’t need to rely on seigniorage. Second, central banks often have enough independence to resist political pressure for more expansionary monetary policy. Third, and most important, policymakers in all parts of government know that inflation is a poor solution to fiscal problems.

Debt and the Political Process

Fiscal policy is made not by angels but by an imperfect political process. Some economists worry that the possibility of financing government spending by issuing debt makes that political process all the worse.

This idea has a long history. Nineteenth-century economist Knut Wicksell claimed that if the benefit of some type of government spending exceeded its cost, it should be possible to finance that spending in a way that would receive unanimous support from the voters. He concluded that government spending should be undertaken only when support was, in fact, nearly unanimous. In the case of debt finance, however, Wicksell was concerned that “the interests [of future taxpayers] are not represented at all or are represented inadequately in the tax-approving assembly.”

Many economists have echoed this theme more recently. In their 1977 book *Democracy in Deficit*, James Buchanan and Richard Wagner argued for a balanced-budget rule for fiscal policy on the grounds that it “will have the effect of bringing the real costs of public outlays to the awareness of decision makers; it will tend to dispel the illusory ‘something for nothing’ aspects of fiscal choice.” Similarly, Martin Feldstein (once an economic adviser to Ronald Reagan and a long-time critic of budget deficits) argues that “only the ‘hard budget constraint’ of having to balance the budget” can force politicians to judge whether spending’s “benefits really justify its costs.”

These arguments have led some economists to favor a constitutional amendment that would require Congress to pass a balanced budget. Often these proposals have escape clauses for times of national emergency, such as wars and depressions, when a budget deficit is a reasonable policy response. Some critics of these proposals argue that, even with the escape clauses, such a constitutional amendment would tie the hands of policymakers too severely. Others claim that Congress would easily evade the balanced-budget requirement with accounting tricks. As this discussion makes clear, the debate over the desirability of a balanced-budget amendment is as much political as economic.

International Dimensions

Government debt may affect a nation’s role in the world economy. As we first saw in Chapter 5, when a government budget deficit reduces national saving, it often leads to a trade deficit, which in turn is financed by borrowing from abroad. For instance, many observers have blamed U.S. fiscal policy for the recent switch of the United States from a major creditor in the world economy to a major debtor. This link between the budget deficit and the trade deficit leads to two further effects of government debt.

First, high levels of government debt may increase the risk that an economy will experience capital flight—an abrupt decline in the demand for a country’s assets in world financial markets. International investors are aware that a government can always deal with its debt simply by defaulting. This approach was used as far back as 1335, when England’s King Edward III defaulted

on his debt to Italian bankers. More recently, several Latin American countries defaulted on their debts in the 1980s, and Russia did the same in 1998. The higher the level of the government debt, the greater the temptation of default. Thus, as government debt increases, international investors may come to fear default and curtail their lending. If this loss of confidence occurs suddenly, the result could be the classic symptoms of capital flight: a collapse in the value of the currency and an increase in interest rates. As we discussed in Chapter 12, this is precisely what happened to Mexico in the early 1990s when default appeared likely.

Second, high levels of government debt financed by foreign borrowing may reduce a nation's political clout in world affairs. This fear was emphasized by economist Ben Friedman in his 1988 book *Day of Reckoning*. He wrote, "World power and influence have historically accrued to creditor countries. It is not coincidental that America emerged as a world power simultaneously with our transition from a debtor nation . . . to a creditor supplying investment capital to the rest of the world." Friedman suggests that if the United States continues to run large trade deficits, it will eventually lose some of its international influence. So far, the record has not been kind to this hypothesis: the United States has run another decade of trade deficits and remains a leading superpower. But perhaps other events—such as the collapse of the Soviet Union—offset the fall in political clout that the United States would have experienced from its increased indebtedness.

CASE STUDY

The Benefits of Indexed Bonds

In 1997, the U.S. Treasury Department started to issue bonds that pay a return based on the consumer price index. These bonds pay a low interest rate of about 3.5 percent, so a \$1,000 bond pays only \$35 per year in interest. But that interest payment grows with the overall price level as measured by the CPI. In addition, when the \$1,000 of principal is repaid, that amount is also adjusted for changes in the CPI. The 3.5 percent, therefore, is a real interest rate. No longer do professors of macroeconomics need to define the real interest rate as an abstract construct. They can open the *New York Times*, point to the credit report, and say, "Look here, this is a nominal interest rate, and this is a real interest rate." (Professors in the United Kingdom and several other countries have long enjoyed this luxury because indexed bonds have been trading in other countries for years.)

Of course, making macroeconomics easier to teach was not the reason that the Treasury chose to index some of the government debt. That was just a positive externality. Its goal was to introduce a new type of government bond that should benefit bondholder and taxpayer alike. These bonds are a win-win proposition because they insulate both sides of the transaction from inflation risk. Bondholders should care about the real interest rate they earn, and taxpayers should care about the real interest rate they pay. When government bonds are

specified in nominal terms, both sides take on risk that is neither productive nor necessary. The new indexed bonds eliminate this inflation risk.

In addition, the new bonds have three other benefits:

First, the bonds may encourage the private sector to begin issuing its own indexed securities. Financial innovation is, to some extent, a public good. Once an innovation has been introduced into the market, the idea is nonexcludable (people cannot be prevented from using it) and nonrival (one person's use of the idea does not diminish other people's use of it). Just as a free market will not adequately supply the public goods of national defense and basic research, it will not adequately supply financial innovation. The Treasury's new bonds can be viewed as a remedy for that market failure.

Second, the bonds reduce the government's incentive to produce surprise inflation. After the large budget deficits of the 1980s and 1990s, the U.S. government is now a substantial debtor, and its debts are specified almost entirely in dollar terms. What is unique about the federal government, in contrast to most debtors, is that it can print the money it needs. The greater the government's nominal debts, the more incentive the government has to inflate away its debt. The Treasury's switch toward indexed debt reduces this potentially problematic incentive.

Third, the bonds provide data that might be useful for monetary policy. Many macroeconomic theories point to expected inflation as a key variable to explain the relationship between inflation and unemployment. But what is expected inflation? One way to measure it is to survey private forecasters. Another way is to look at the difference between the yield on nominal bonds and the yield on real bonds.

In the past, economists have proposed a variety of rules that could be used to conduct monetary policy, as we discussed in the preceding chapter. The new indexed bonds expand the number of possible rules. Here is one idea: the Fed announces a target for the inflation rate. Then, every day, the Fed measures expected inflation as the spread between the yield on nominal debt and the yield on indexed debt. If expected inflation is above the target, the Fed contracts the money supply. If expected inflation is below the target, the Fed expands the money supply. In this way, the Fed can use the bond market's inflation forecast to ensure that the money supply is growing at the rate needed to keep inflation close to its target.

The Treasury's new indexed bonds, therefore, will likely produce many benefits: less inflation risk, more financial innovation, better government incentives, more informed monetary policy, and easier lives for students and teachers of macroeconomics.⁷

⁷ To read more about indexed bonds, see John Y. Campbell and Robert J. Shiller, "A Scorecard for Indexed Government Debt," *NBER Macroeconomics Annual* (1996): 155–197; and David W. Wilcox, "Policy Watch: The Introduction of Indexed Government Debt in the United States," *The Journal of Economic Perspectives* 12 (Winter 1998): 219–227.

15-6 Conclusion

Fiscal policy and government debt are central in the U.S. political debate. When Bill Clinton became president in 1993, he made reducing the budget deficit a high priority of his administration. When the Republicans took control of Congress in 1995, they pushed for even faster deficit reduction than Clinton had advocated. These efforts together with some good luck turned the federal government budget from deficit to surplus by the late 1990s. When George W. Bush moved into the White House in 2001, the policy debate was over how quickly the government should pay off its debts.

This chapter discussed some of the economic issues that lie behind these policy decisions. As we have seen, economists are not in complete agreement about the measurement or effects of government indebtedness. Given the profound importance of this topic, there seems little doubt that the debates will continue in the years to come.

Summary

1. The current debt of the U.S. federal government is of moderate size compared to the debt of other countries or compared to the debt that the United States has had throughout its own history. The 1980s and early 1990s were unusual in that the ratio of debt to GDP increased during a period of peace and prosperity. Since 1995, the debt–GDP ratio has declined substantially.
2. Standard measures of the budget deficit are imperfect measures of fiscal policy because they do not correct for the effects of inflation, do not offset changes in government liabilities with changes in government assets, omit some liabilities altogether, and do not correct for the effects of the business cycle.
3. According to the traditional view of government debt, a debt-financed tax cut stimulates consumer spending and lowers national saving. This increase in consumer spending leads to greater aggregate demand and higher income in the short run, but it leads to a lower capital stock and lower income in the long run.
4. According to the Ricardian view of government debt, a debt-financed tax cut does not stimulate consumer spending because it does not raise consumers' overall resources—it merely reschedules taxes from the present to the future. The debate between the traditional and Ricardian views of government debt is ultimately a debate over how consumers behave. Are consumers rational or shortsighted? Do they face binding borrowing constraints? Are they economically linked to future generations through altruistic bequests? Economists' views of government debt hinge on their answers to these questions.

5. Most economists oppose a strict rule requiring a balanced budget. A budget deficit can sometimes be justified on the basis of short-run stabilization, tax smoothing, or intergenerational redistribution of the tax burden.
6. Government debt can potentially have other effects. Large government debt or budget deficits may encourage excessive monetary expansion and, therefore, lead to greater inflation. The possibility of running budget deficits may encourage politicians to unduly burden future generations when setting government spending and taxes. A high level of government debt may risk capital flight and diminish a nation's influence around the world. Economists differ in which of these effects they consider most important.

KEY CONCEPTS

Capital budgeting

Cyclically adjusted budget deficit

Ricardian equivalence

QUESTIONS FOR REVIEW

1. What was unusual about U.S. fiscal policy from 1980 to 1995?
2. Why do many economists project increasing budget deficits and government debt over the next several decades?
3. Describe four problems affecting measurement of the government budget deficit.
4. According to the traditional view of government debt, how does a debt-financed tax cut affect public saving, private saving, and national saving?
5. According to the Ricardian view of government debt, how does a debt-financed tax cut affect public saving, private saving, and national saving?
6. Do you believe the traditional or the Ricardian view of government debt? Why?
7. Give three reasons why a budget deficit might be a good policy choice.
8. Why might the level of government debt affect the government's incentives regarding money creation?

PROBLEMS AND APPLICATIONS

1. On April 1, 1996, Taco Bell, the fast-food chain, ran a full-page ad in the *New York Times* with this news: "In an effort to help the national debt, Taco Bell is pleased to announce that we have agreed to purchase the Liberty Bell, one of our country's most historic treasures. It will now be called the *Taco Liberty Bell* and will still be accessible to the American public for viewing. We hope our move will prompt other corporations to take similar action to do their part to reduce the country's debt." Would such actions by U.S. corporations actually reduce the national debt as it is now measured? How would your answer change if the U.S. government adopted capital budgeting? Do you think these actions represent a true reduction in the government's indebtedness? Do you think Taco Bell was serious about this plan? (*Hint:* Note the date.)
2. Draft a letter to the senator described in Section 15-3, explaining and evaluating the Ricardian view of government debt.
3. The Social Security system levies a tax on workers and pays benefits to the elderly. Suppose that

Congress increases both the tax and the benefits. For simplicity, assume that the Congress announces that the increases will last for one year only.

- a. How do you suppose this change would affect the economy? (*Hint:* Think about the marginal propensities to consume of the young and the old.)
 - b. Does your answer depend on whether generations are altruistically linked?
4. Evaluate the usefulness of generational accounting from the perspective of someone who believes that generations are altruistically linked. Now evaluate the usefulness of generational accounting from the perspective of someone who believes that many consumers face binding borrowing constraints.
 5. The *cyclically adjusted budget deficit* is the budget deficit corrected for the effects of the business

cycle. In other words, it is the budget deficit that the government would be running if unemployment were at the natural rate. (It is also called the *full-employment budget deficit*.) Some economists have proposed the rule that the cyclically adjusted budget deficit always be balanced. Compare this proposal to a strict balanced-budget rule. Which is preferable? What problems do you see with the rule requiring a balanced cyclically adjusted budget?

6. Using the library or the Internet, find some recent projections for the future path of the U.S. government debt as a percentage of GDP. What assumptions are made about government spending, taxes, and economic growth? Do you think these assumptions are reasonable? If the U.S. experiences a productivity slowdown, how will reality differ from this projection? (*Hint:* A good place to look is www.cbo.gov.)

part VI

**More on the
Microeconomics
Behind Macroeconomics**

Consumption

Consumption is the sole end and purpose of all production.

— Adam Smith

How do households decide how much of their income to consume today and how much to save for the future? This is a microeconomic question because it addresses the behavior of individual decisionmakers. Yet its answer has macroeconomic consequences. As we have seen in previous chapters, households' consumption decisions affect the way the economy as a whole behaves both in the long run and in the short run.

The consumption decision is crucial for long-run analysis because of its role in economic growth. The Solow growth model of Chapters 7 and 8 shows that the saving rate is a key determinant of the steady-state capital stock and thus of the level of economic well-being. The saving rate measures how much of its income the present generation is putting aside for its own future and for future generations.

The consumption decision is crucial for short-run analysis because of its role in determining aggregate demand. Consumption is two-thirds of GDP, so fluctuations in consumption are a key element of booms and recessions. The *IS-LM* model of Chapters 10 and 11 shows that changes in consumers' spending plans can be a source of shocks to the economy, and that the marginal propensity to consume is a determinant of the fiscal-policy multipliers.

In previous chapters we explained consumption with a function that relates consumption to disposable income: $C = C(Y - T)$. This approximation allowed us to develop simple models for long-run and short-run analysis, but it is too simple to provide a complete explanation of consumer behavior. In this chapter we examine the consumption function in greater detail and develop a more thorough explanation of what determines aggregate consumption.

Since macroeconomics began as a field of study, many economists have written about the theory of consumer behavior and suggested alternative ways of interpreting the data on consumption and income. This chapter presents the views of six prominent economists to show the diverse approaches to explaining consumption.

16-1 John Maynard Keynes and the Consumption Function

We begin our study of consumption with John Maynard Keynes's *General Theory*, which was published in 1936. Keynes made the consumption function central to his theory of economic fluctuations, and it has played a key role in macroeconomic analysis ever since. Let's consider what Keynes thought about the consumption function, and then see what puzzles arose when his ideas were confronted with the data.

Keynes's Conjectures


Today, economists who study consumption rely on sophisticated techniques of data analysis. With the help of computers, they analyze aggregate data on the behavior of the overall economy from the national income accounts and detailed data on the behavior of individual households from surveys. Because Keynes wrote in the 1930s, however, he had neither the advantage of these data nor the computers necessary to analyze such large data sets. Instead of relying on statistical analysis, Keynes made conjectures about the consumption function based on introspection and casual observation.

First and most important, Keynes conjectured that the **marginal propensity to consume**—the amount consumed out of an additional dollar of income—is between zero and one. He wrote that the “fundamental psychological law, upon which we are entitled to depend with great confidence, . . . is that men are disposed, as a rule and on the average, to increase their consumption as their income increases, but not by as much as the increase in their income.” That is, when a person earns an extra dollar, he typically spends some of it and saves some of it. As we saw in Chapter 10 when we developed the Keynesian cross, the marginal propensity to consume was crucial to Keynes's policy recommendations for how to reduce widespread unemployment. The power of fiscal policy to influence the economy—as expressed by the fiscal-policy multipliers—arises from the feedback between income and consumption.

Second, Keynes posited that the ratio of consumption to income, called the **average propensity to consume**, falls as income rises. He believed that saving was a luxury, so he expected the rich to save a higher proportion of their income than the poor. Although not essential for Keynes's own analysis, the postulate that the average propensity to consume falls as income rises became a central part of early Keynesian economics.

Third, Keynes thought that income is the primary determinant of consumption and that the interest rate does not have an important role. This conjecture stood in stark contrast to the beliefs of the classical economists who preceded him. The classical economists held that a higher interest rate encourages saving and discourages consumption. Keynes admitted that the interest rate could influence consumption as a matter of theory. Yet he wrote that “the main conclusion suggested

figure 16-1



Note: The marginal propensity to consume, MPC , is the slope of the consumption function. The average propensity to consume, $APC = C/Y$, equals the slope of a line drawn from the origin to a point on the consumption function.

The Keynesian Consumption Function This figure graphs a consumption function with the three properties that Keynes conjectured. First, the marginal propensity to consume c is between zero and one. Second, the average propensity to consume falls as income rises. Third, consumption is determined by current income.

by experience, I think, is that the short-period influence of the rate of interest on individual spending out of a given income is secondary and relatively unimportant.”

On the basis of these three conjectures, the Keynesian consumption function is often written as

$$C = \bar{C} + cY, \quad \bar{C} > 0, \quad 0 < c < 1,$$

where C is consumption, Y is disposable income, $\bar{C}$ is a constant, and c is the marginal propensity to consume. This consumption function, shown in Figure 16-1, is graphed as a straight line.

Notice that this consumption function exhibits the three properties that Keynes posited. It satisfies Keynes’s first property because the marginal propensity to consume c is between zero and one, so that higher income leads to higher consumption and also to higher saving. This consumption function satisfies Keynes’s second property because the average propensity to consume APC is

$$APC = C/Y = \bar{C}/Y + c.$$

As Y rises, $\bar{C}/Y$ falls, and so the average propensity to consume C/Y falls. And finally, this consumption function satisfies Keynes’s third property because the interest rate is not included in this equation as a determinant of consumption.

The Early Empirical Successes

Soon after Keynes proposed the consumption function, economists began collecting and examining data to test his conjectures. The earliest studies indicated that the Keynesian consumption function is a good approximation of how consumers behave.

In some of these studies, researchers surveyed households and collected data on consumption and income. They found that households with higher income consumed more, which confirms that the marginal propensity to consume is greater than zero. They also found that households with higher income saved more, which confirms that the marginal propensity to consume is less than one. In addition, these researchers found that higher-income households saved a larger fraction of their income, which confirms that the average propensity to consume falls as income rises. Thus, these data verified Keynes's conjectures about the marginal and average propensities to consume.

In other studies, researchers examined aggregate data on consumption and income for the period between the two world wars. These data also supported the Keynesian consumption function. In years when income was unusually low, such as during the depths of the Great Depression, both consumption and saving were low, indicating that the marginal propensity to consume is between zero and one. In addition, during those years of low income, the ratio of consumption to income was high, confirming Keynes's second conjecture. Finally, because the correlation between income and consumption was so strong, no other variable appeared to be important for explaining consumption. Thus, the data also confirmed Keynes's third conjecture that income is the primary determinant of how much people choose to consume.

Secular Stagnation, Simon Kuznets, and the Consumption Puzzle

Although the Keynesian consumption function met with early successes, two anomalies soon arose. Both concern Keynes's conjecture that the average propensity to consume falls as income rises.

The first anomaly became apparent after some economists made a dire—and, it turned out, erroneous—prediction during World War II. On the basis of the Keynesian consumption function, these economists reasoned that as incomes in the economy grew over time, households would consume a smaller and smaller fraction of their incomes. They feared that there might not be enough profitable investment projects to absorb all this saving. If so, the low consumption would lead to an inadequate demand for goods and services, resulting in a depression once the wartime demand from the government ceased. In other words, on the basis of the Keynesian consumption function, these economists predicted that the economy would experience what they called *secular stagnation*—a long depression of indefinite duration—unless fiscal policy was used to expand aggregate demand.

Fortunately for the economy, but unfortunately for the Keynesian consumption function, the end of World War II did not throw the country into another depression. Although incomes were much higher after the war than before, these higher incomes did not lead to large increases in the rate of saving. Keynes's conjecture that the average propensity to consume would fall as income rose appeared not to hold.

The second anomaly arose when economist Simon Kuznets constructed new aggregate data on consumption and income dating back to 1869. Kuznets assembled


these data in the 1940s and would later receive the Nobel Prize for this work. He discovered that the ratio of consumption to income was remarkably stable from decade to decade, despite large increases in income over the period he studied. Again, Keynes's conjecture that the average propensity to consume would fall as income rose appeared not to hold.

The failure of the secular-stagnation hypothesis and the findings of Kuznets both indicated that the average propensity to consume is fairly constant over long periods of time. This fact presented a puzzle that motivated much of the subsequent work on consumption. Economists wanted to know why some studies confirmed Keynes's conjectures and others refuted them. That is, why did Keynes's conjectures hold up well in the studies of household data and in the studies of short time-series, but fail when long time-series were examined?

Figure 16-2 illustrates the puzzle. The evidence suggested that there were two consumption functions. For the household data or for the short time-series, the Keynesian consumption function appeared to work well. Yet for the long time-series, the consumption function appeared to have a constant average propensity to consume. In Figure 16-2, these two relationships between consumption and income are called the short-run and long-run consumption functions. Economists needed to explain how these two consumption functions could be consistent with each other.

In the 1950s, Franco Modigliani and Milton Friedman each proposed explanations of these seemingly contradictory findings. Both economists later won

figure 16-2



The Consumption Puzzle

Studies of household data and short time-series found a relationship between consumption and income similar to the one Keynes conjectured. In the figure, this relationship is called the short-run consumption function. But studies of long time-series found that the average propensity to consume did not vary systematically with income. This relationship is called the long-run consumption function. Notice that the short-run consumption function has a falling average propensity to consume, whereas the long-run consumption function has a constant average propensity to consume.

Nobel Prizes, in part because of their work on consumption. But before we see how Modigliani and Friedman tried to solve the consumption puzzle, we must discuss Irving Fisher's contribution to consumption theory. Both Modigliani's life-cycle hypothesis and Friedman's permanent-income hypothesis rely on the theory of consumer behavior proposed much earlier by Irving Fisher.

16-2 Irving Fisher and Intertemporal Choice

The consumption function introduced by Keynes relates current consumption to current income. This relationship, however, is incomplete at best. When people decide how much to consume and how much to save, they consider both the present and the future. The more consumption they enjoy today, the less they will be able to enjoy tomorrow. In making this tradeoff, households must look ahead to the income they expect to receive in the future and to the consumption of goods and services they hope to be able to afford.

The economist Irving Fisher developed the model with which economists analyze how rational, forward-looking consumers make intertemporal choices—that is, choices involving different periods of time. Fisher's model illuminates the constraints consumers face, the preferences they have, and how these constraints and preferences together determine their choices about consumption and saving.

The Intertemporal Budget Constraint

Most people would prefer to increase the quantity or quality of the goods and services they consume—to wear nicer clothes, eat at better restaurants, or see more movies. The reason people consume less than they desire is that their consumption is constrained by their income. In other words, consumers face a limit on how much they can spend, called a *budget constraint*. When they are deciding how much to consume today versus how much to save for the future, they face an **intertemporal budget constraint**, which measures the total resources available for consumption today and in the future. Our first step in developing Fisher's model is to examine this constraint in some detail.

To keep things simple, we examine the decision facing a consumer who lives for two periods. Period one represents the consumer's youth, and period two represents the consumer's old age. The consumer earns income Y_1 and consumes C_1 in period one, and earns income Y_2 and consumes C_2 in period two. (All variables are real—that is, adjusted for inflation.) Because the consumer has the opportunity to borrow and save, consumption in any single period can be either greater or less than income in that period.

Consider how the consumer's income in the two periods constrains consumption in the two periods. In the first period, saving equals income minus consumption. That is,

$$S = Y_1 - C_1,$$

where S is saving. In the second period, consumption equals the accumulated saving, including the interest earned on that saving, plus second-period income. That is,

$$C_2 = (1 + r)S + Y_2,$$

where r is the real interest rate. For example, if the interest rate is 5 percent, then for every \$1 of saving in period one, the consumer enjoys an extra \$1.05 of consumption in period two. Because there is no third period, the consumer does not save in the second period.

Note that the variable S can represent either saving or borrowing and that these equations hold in both cases. If first-period consumption is less than first-period income, the consumer is saving, and S is greater than zero. If first-period consumption exceeds first-period income, the consumer is borrowing, and S is less than zero. For simplicity, we assume that the interest rate for borrowing is the same as the interest rate for saving.

To derive the consumer's budget constraint, combine the two preceding equations. Substitute the first equation for S into the second equation to obtain

$$C_2 = (1 + r)(Y_1 - C_1) + Y_2.$$

To make the equation easier to interpret, we must rearrange terms. To place all the consumption terms together, bring $(1 + r)C_1$ from the right-hand side to the left-hand side of the equation to obtain

$$(1 + r)C_1 + C_2 = (1 + r)Y_1 + Y_2.$$

Now divide both sides by $1 + r$ to obtain

$$C_1 + \frac{C_2}{1 + r} = Y_1 + \frac{Y_2}{1 + r}.$$

This equation relates consumption in the two periods to income in the two periods. It is the standard way of expressing the consumer's intertemporal budget constraint.

The consumer's budget constraint is easily interpreted. If the interest rate is zero, the budget constraint shows that total consumption in the two periods equals total income in the two periods. In the usual case in which the interest rate is greater than zero, future consumption and future income are discounted by a factor $1 + r$. This **discounting** arises from the interest earned on savings. In essence, because the consumer earns interest on current income that is saved, future income is worth less than current income. Similarly, because future consumption is paid for out of savings that have earned interest, future consumption costs less than current consumption. The factor $1/(1 + r)$ is the price of second-period consumption measured in terms of first-period consumption: it is the amount of first-period consumption that the consumer must forgo to obtain 1 unit of second-period consumption.

Figure 16-3 graphs the consumer's budget constraint. Three points are marked on this figure. At point A, the consumer consumes exactly his income in each

FYI

Present Value, or Why a \$1,000,000 Prize Is Worth Only \$623,000

The use of discounting in the consumer's budget constraint illustrates an important fact of economic life: a dollar in the future is less valuable than a dollar today. This is true because a dollar today can be deposited in an interest-bearing bank account and produce more than one dollar in the future. If the interest rate is 5 percent, for instance, then a dollar today can be turned to \$1.05 dollars next year, \$1.1025 in two years, \$1.1576 in three years, . . . , or \$2.65 in 20 years.

Economists use a concept called *present value* to compare dollar amounts from different times. The present value of any amount in the future is the amount that would be needed today, given available interest rates, to produce that future amount. Thus, if you are going to be paid X dollars in T years and the interest rate is r , then the present value of that payment is


$$\text{Present Value} = X/(1 + r)^T.$$

In light of this definition, we can see a new interpretation of the consumer's budget constraint in our two-period consumption problem. The intertemporal budget constraint states that the present value of consumption must equal the present value of income.

The concept of present value has many applications. Suppose, for instance, that you won a million-dollar lottery. Such prizes are usually paid out over time—say, \$50,000 a year for 20 years. What is the present value of such a delayed prize? By applying the preceding formula for each of the 20 payments and adding up the result, we learn that the million-dollar prize, discounted at an interest rate of 5 percent, has a present value of only \$623,000. (If the prize were paid out as a dollar a year for a million years, the present value would be a mere \$20!) Sometimes a million dollars isn't all it's cracked up to be.

period ($C_1 = Y_1$ and $C_2 = Y_2$), so there is neither saving nor borrowing between the two periods. At point B, the consumer consumes nothing in the first period ($C_1 = 0$) and saves all income, so second-period consumption C_2 is $(1 + r)Y_1 + Y_2$. At point C, the consumer plans to consume nothing in the second period ($C_2 = 0$) and borrows as much as possible against second-period income, so first-period

figure 16-3



The Consumer's Budget Constraint This figure shows the combinations of first-period and second-period consumption the consumer can choose. If he chooses points between A and B, he consumes less than his income in the first period and saves the rest for the second period. If he chooses points between A and C, he consumes more than his income in the first period and borrows to make up the difference.

consumption C_1 is $Y_1 + Y_2/(1 + r)$. Of course, these are only three of the many combinations of first- and second-period consumption that the consumer can afford: all the points on the line from B to C are available to the consumer.

Consumer Preferences

The consumer's preferences regarding consumption in the two periods can be represented by **indifference curves**. An indifference curve shows the combinations of first-period and second-period consumption that make the consumer equally happy.


Figure 16-4 shows two of the consumer's many indifference curves. The consumer is indifferent among combinations W, X, and Y, because they are all on the same curve. Not surprisingly, if the consumer's first-period consumption is reduced, say from point W to point X, second-period consumption must increase to keep him equally happy. If first-period consumption is reduced again, from point X to point Y, the amount of extra second-period consumption he requires for compensation is greater.

The slope at any point on the indifference curve shows how much second-period consumption the consumer requires in order to be compensated for a 1-unit reduction in first-period consumption. This slope is the **marginal rate of substitution** between first-period consumption and second-period consumption. It tells us the rate at which the consumer is willing to substitute second-period consumption for first-period consumption.

Notice that the indifference curves in Figure 16-4 are not straight lines and, as a result, the marginal rate of substitution depends on the levels of consumption in the two periods. When first-period consumption is high and second-period consumption is low, as at point W, the marginal rate of substitution is low: the consumer requires only a little extra second-period consumption to give up

figure 16-4

Second-period consumption, C_2



The Consumer's Preferences

Indifference curves represent the consumer's preferences over first-period and second-period consumption. An indifference curve gives the combinations of consumption in the two periods that make the consumer equally happy. This figure shows two of many indifference curves. Higher indifference curves such as IC_2 are preferred to lower curves such as IC_1 . The consumer is equally happy at points W, X, and Y, but prefers point Z to points W, X, or Y.

1 unit of first-period consumption. When first-period consumption is low and second-period consumption is high, as at point Y, the marginal rate of substitution is high: the consumer requires much additional second-period consumption to give up 1 unit of first-period consumption.

The consumer is equally happy at all points on a given indifference curve, but he prefers some indifference curves to others. Because he prefers more consumption to less, he prefers higher indifference curves to lower ones. In Figure 16-4, the consumer prefers the points on curve IC_2 to the points on curve IC_1 .


The set of indifference curves gives a complete ranking of the consumer's preferences. It tells us that the consumer prefers point Z to point W, but that may be obvious because point Z has more consumption in both periods. Yet compare point Z and point Y: point Z has more consumption in period one and less in period two. Which is preferred, Z or Y? Because Z is on a higher indifference curve than Y, we know that the consumer prefers point Z to point Y. Hence, we can use the set of indifference curves to rank any combinations of first-period and second-period consumption.

Optimization

Having discussed the consumer's budget constraint and preferences, we can consider the decision about how much to consume. The consumer would like to end up with the best possible combination of consumption in the two periods—that is, on the highest possible indifference curve. But the budget constraint requires that the consumer also end up on or below the budget line, because the budget line measures the total resources available to him.

Figure 16-5 shows that many indifference curves cross the budget line. The highest indifference curve that the consumer can obtain without violating the budget constraint is the indifference curve that just barely touches the budget

figure 16-5



The Consumer's Optimum

The consumer achieves his highest level of satisfaction by choosing the point on the budget constraint that is on the highest indifference curve. At the optimum, the indifference curve is tangent to the budget constraint.

line, which is curve IC_3 in the figure. The point at which the curve and line touch—point O for “optimum”—is the best combination of consumption in the two periods that the consumer can afford.

Notice that, at the optimum, the slope of the indifference curve equals the slope of the budget line. The indifference curve is *tangent* to the budget line. The slope of the indifference curve is the marginal rate of substitution MRS , and the slope of the budget line is 1 plus the real interest rate. We conclude that at point O,

$$MRS = 1 + r.$$

The consumer chooses consumption in the two periods so that the marginal rate of substitution equals 1 plus the real interest rate.


How Changes in Income Affect Consumption

Now that we have seen how the consumer makes the consumption decision, let’s examine how consumption responds to an increase in income. An increase in either Y_1 or Y_2 shifts the budget constraint outward, as in Figure 16-6. The higher budget constraint allows the consumer to choose a better combination of first- and second-period consumption—that is, the consumer can now reach a higher indifference curve.

In Figure 16-6, the consumer responds to the shift in his budget constraint by choosing more consumption in both periods. Although it is not implied by the logic of the model alone, this situation is the most usual. If a consumer wants more of a good when his or her income rises, economists call it a **normal good**. The indifference curves in Figure 16-6 are drawn under the assumption that consumption in period one and consumption in period two are both normal goods.

The key conclusion from Figure 16-6 is that regardless of whether the increase in income occurs in the first period or the second period, the consumer spreads it over consumption in both periods. This behavior is sometimes called

figure 16-6



An Increase in Income An increase in either first-period income or second-period income shifts the budget constraint outward. If consumption in period one and consumption in period two are both normal goods, this increase in income raises consumption in both periods.

consumption smoothing. Because the consumer can borrow and lend between periods, the timing of the income is irrelevant to how much is consumed today (except, of course, that future income is discounted by the interest rate). The lesson of this analysis is that consumption depends on the present value of current and future income—that is, on

$$\text{Present Value of Income} = Y_1 + \frac{Y_2}{1 + r}.$$

Notice that this conclusion is quite different from that reached by Keynes. *Keynes posited that a person's current consumption depends largely on his current income. Fisher's model says, instead, that consumption is based on the resources the consumer expects over his lifetime.*

How Changes in the Real Interest Rate Affect Consumption

Let's now use Fisher's model to consider how a change in the real interest rate alters the consumer's choices. There are two cases to consider: the case in which the consumer is initially saving and the case in which he is initially borrowing. Here we discuss the saving case; Problem 1 at the end of the chapter asks you to analyze the borrowing case.

Figure 16-7 shows that an increase in the real interest rate rotates the consumer's budget line around the point (Y_1, Y_2) and, thereby, alters the amount of consumption he chooses in both periods. Here, the consumer moves from point A to point B. You can see that for the indifference curves drawn in this figure first-period consumption falls and second-period consumption rises.


Economists decompose the impact of an increase in the real interest rate on consumption into two effects: an **income effect** and a **substitution effect**. Textbooks in microeconomics discuss these effects in detail. We summarize them briefly here.

The income effect is the change in consumption that results from the movement to a higher indifference curve. Because the consumer is a saver rather than a borrower (as indicated by the fact that first-period consumption is less than first-period income), the increase in the interest rate makes him better off (as reflected by the movement to a higher indifference curve). If consumption in period one and consumption in period two are both normal goods, the consumer will want to spread this improvement in his welfare over both periods. This income effect tends to make the consumer want more consumption in both periods.

The substitution effect is the change in consumption that results from the change in the relative price of consumption in the two periods. In particular, consumption in period two becomes less expensive relative to consumption in period one when the interest rate rises. That is, because the real interest rate earned on saving is higher, the consumer must now give up less first-period consumption to obtain an extra unit of second-period consumption. This substitution effect tends to make the consumer choose more consumption in period two and less consumption in period one.

The consumer's choice depends on both the income effect and the substitution effect. Both effects act to increase the amount of second-period consumption;

figure 16-7

**An Increase in the Interest Rate**

An increase in the interest rate rotates the budget constraint around the point (Y_1, Y_2) . In this figure, the higher interest rate reduces first-period consumption by ΔC_1 and raises second-period consumption by ΔC_2 .

hence, we can confidently conclude that an increase in the real interest rate raises second-period consumption. But the two effects have opposite impacts on first-period consumption, so the increase in the interest rate could either raise or lower it. *Hence, depending on the relative size of the income and substitution effects, an increase in the interest rate could either stimulate or depress saving.*

Constraints on Borrowing

Fisher's model assumes that the consumer can borrow as well as save. The ability to borrow allows current consumption to exceed current income. In essence, when the consumer borrows, he consumes some of his future income today. Yet for many people such borrowing is impossible. For example, a student wishing to enjoy spring break in Florida would probably be unable to finance this vacation with a bank loan. Let's examine how Fisher's analysis changes if the consumer cannot borrow.

The inability to borrow prevents current consumption from exceeding current income. A constraint on borrowing can therefore be expressed as


$$C_1 \leq Y_1.$$

This inequality states that consumption in period one must be less than or equal to income in period one. This additional constraint on the consumer is called a **borrowing constraint** or, sometimes, a *liquidity constraint*.

Figure 16-8 shows how this borrowing constraint restricts the consumer's set of choices. The consumer's choice must satisfy both the intertemporal budget constraint and the borrowing constraint. The shaded area represents the combinations of first-period consumption and second-period consumption that satisfy both constraints.

Figure 16-9 shows how this borrowing constraint affects the consumption decision. There are two possibilities. In panel (a), the consumer wishes to consume less in period one than he earns. The borrowing constraint is not binding and, therefore, does not affect consumption. In panel (b), the consumer would like to choose point D, where he consumes more in period one than he earns, but the borrowing constraint prevents this outcome. The best the consumer can do is to consume his first-period income, represented by point E.


The analysis of borrowing constraints leads us to conclude that there are two consumption functions. For some consumers, the borrowing constraint is not binding, and consumption in both periods depends on the present value of lifetime income, $Y_1 + [Y_2/(1 + r)]$. For other consumers, the borrowing constraint binds, and the consumption function is $C_1 = Y_1$ and $C_2 = Y_2$. Hence, for those consumers who would like to borrow but cannot, consumption depends only on current income.



Drawing by Handelsman; © 1985
The New Yorker Magazine, Inc.


"What I'd like, basically, is a temporary line of credit just to tide me over the rest of my life."

figure 16-8



A Borrowing Constraint If the consumer cannot borrow, he faces the additional constraint that first-period consumption cannot exceed first-period income. The shaded area represents the combination of first-period and second-period consumption the consumer can choose.

figure 16-9



The Consumer's Optimum With a Borrowing Constraint When the consumer faces a borrowing constraint, there are two possible situations. In panel (a), the consumer chooses first-period consumption to be less than first-period income, so the borrowing constraint is not binding and does not affect consumption in either period. In panel (b), the borrowing constraint is binding. The consumer would like to borrow and choose point D. But because borrowing is not allowed, the best available choice is point E. When the borrowing constraint is binding, first-period consumption equals first-period income.

CASE STUDY

The High Japanese Saving Rate

Japan has one of the world's highest saving rates, and this fact is important for understanding both the long-run and short-run performance of its economy. On the one hand, many economists believe that the high Japanese saving rate is a key to the rapid growth Japan experienced in the decades after World War II. Indeed, the Solow growth model developed in Chapters 7 and 8 shows that the saving rate is a primary determinant of a country's steady-state level of income. On the other hand, some economists have argued that the high Japanese saving rate contributed to Japan's slump during the 1990s. High saving means low consumer spending, which according to the *IS-LM* model of Chapters 10 and 11 translates into low aggregate demand and reduced income.

Why do the Japanese consume a much smaller fraction of their income than Americans? One reason is that it is harder for households to borrow in

Japan. As Fisher's model shows, a household facing a binding borrowing constraint consumes less than it would without the borrowing constraint. Hence, societies in which borrowing constraints are common will tend to have higher rates of saving.

One reason that households often wish to borrow is to buy a home. In the United States, a person can usually buy a home with a down payment of 10 percent. A homebuyer in Japan cannot borrow nearly this much: down payments of 40 percent are common. Moreover, housing prices are very high in Japan, primarily because land prices are high. A Japanese family must save a great deal if it is ever to afford its own home.

Although constraints on borrowing are part of the explanation of high Japanese saving, there are many other differences between Japan and the United States that contribute to the difference in the saving rates. The Japanese tax system encourages saving by taxing capital income very lightly. In addition, cultural differences may lead to differences in consumer preferences regarding present and future consumption. One prominent Japanese economist writes, "The Japanese are simply *different*. They are more risk averse and more patient. If this is true, the long-run implication is that Japan will absorb all the wealth in the world. I refuse to comment on this explanation."¹

16-3 Franco Modigliani and the Life-Cycle Hypothesis

In a series of papers written in the 1950s, Franco Modigliani and his collaborators Albert Ando and Richard Brumberg used Fisher's model of consumer behavior to study the consumption function. One of their goals was to solve the consumption puzzle—that is, to explain the apparently conflicting pieces of evidence that came to light when Keynes's consumption function was brought to the data. According to Fisher's model, consumption depends on a person's lifetime income. Modigliani emphasized that income varies systematically over people's lives and that saving allows consumers to move income from those times in life when income is high to those times when it is low. This interpretation of consumer behavior formed the basis for his **life-cycle hypothesis**.²

¹ Fumio Hayashi, "Why Is Japan's Saving Rate So Apparently High?" *NBER Macroeconomics Annual* (1986): 147–210.

² For references to the large body of work on the life-cycle hypothesis, a good place to start is the lecture Modigliani gave when he won the Nobel Prize. Franco Modigliani, "Life Cycle, Individual Thrift, and the Wealth of Nations," *American Economic Review* 76 (June 1986): 297–313.

The Hypothesis

One important reason that income varies over a person's life is retirement. Most people plan to stop working at about age 65, and they expect their incomes to fall when they retire. Yet they do not want a large drop in their standard of living, as measured by their consumption. To maintain consumption after retirement, people must save during their working years. Let's see what this motive for saving implies for the consumption function.

Consider a consumer who expects to live another T years, has wealth of W , and expects to earn income Y until she retires R years from now. What level of consumption will the consumer choose if she wishes to maintain a smooth level of consumption over her life?

The consumer's lifetime resources are composed of initial wealth W and lifetime earnings of $R \times Y$. (For simplicity, we are assuming an interest rate of zero; if the interest rate were greater than zero, we would need to take account of interest earned on savings as well.) The consumer can divide up her lifetime resources among her T remaining years of life. We assume that she wishes to achieve the smoothest possible path of consumption over her lifetime. Therefore, she divides this total of $W + RY$ equally among the T years and each year consumes

$$C = (W + RY)/T.$$

We can write this person's consumption function as

$$C = (1/T)W + (R/T)Y.$$

For example, if the consumer expects to live for 50 more years and work for 30 of them, then $T = 50$ and $R = 30$, so her consumption function is

$$C = 0.02W + 0.6Y.$$

This equation says that consumption depends on both income and wealth. An extra \$1 of income per year raises consumption by \$0.60 per year, and an extra \$1 of wealth raises consumption by \$0.02 per year.

If every individual in the economy plans consumption like this, then the aggregate consumption function is much the same as the individual one. In particular, aggregate consumption depends on both wealth and income. That is, the economy's consumption function is


$$C = \alpha W + \beta Y,$$

where the parameter α is the marginal propensity to consume out of wealth, and the parameter β is the marginal propensity to consume out of income.

Implications

Figure 16-10 graphs the relationship between consumption and income predicted by the life-cycle model. For any given level of wealth W , the model yields a conventional consumption function similar to the one shown in Figure 16-1.

figure 16-10



The Life-Cycle Consumption Function The life-cycle model says that consumption depends on wealth as well as income. In other words, the intercept of the consumption function depends on wealth.

Notice, however, that the intercept of the consumption function, which shows what would happen to consumption if income ever fell to zero, is not a fixed value, as it is in Figure 16-1. Instead, the intercept here is αW and, thus, depends on the level of wealth.

This life-cycle model of consumer behavior can solve the consumption puzzle. According to the life-cycle consumption function, the average propensity to consume is


$$C/Y = \alpha(W/Y) + \beta.$$

Because wealth does not vary proportionately with income from person to person or from year to year, we should find that high income corresponds to a low average propensity to consume when looking at data across individuals or over short periods of time. But, over long periods of time, wealth and income grow together, resulting in a constant ratio W/Y and thus a constant average propensity to consume.

To make the same point somewhat differently, consider how the consumption function changes over time. As Figure 16-10 shows, for any given level of wealth, the life-cycle consumption function looks like the one Keynes suggested. But this function holds only in the short run when wealth is constant. In the long run, as wealth increases, the consumption function shifts upward, as in Figure 16-11. This upward shift prevents the average propensity to consume from falling as income increases. In this way, Modigliani resolved the consumption puzzle posed by Simon Kuznets's data.

The life-cycle model makes many other predictions as well. Most important, it predicts that saving varies over a person's lifetime. If a person begins adulthood with no wealth, she will accumulate wealth during her working


figure 16-11



How Changes in Wealth Shift the Consumption Function If consumption depends on wealth, then an increase in wealth shifts the consumption function upward. Thus, the short-run consumption function (which holds wealth constant) will not continue to hold in the long run (as wealth rises over time).

years and then run down her wealth during her retirement years. Figure 16-12 illustrates the consumer's income, consumption, and wealth over her adult life. According to the life-cycle hypothesis, because people want to smooth consumption over their lives, the young who are working save, while the old who are retired dissave.

figure 16-12



Consumption, Income, and Wealth Over the Life Cycle If the consumer smooths consumption over her life (as indicated by the horizontal consumption line), she will save and accumulate wealth during her working years and then dissave and run down her wealth during retirement.

CASE STUDY

The Consumption and Saving of the Elderly

Many economists have studied the consumption and saving of the elderly. Their findings present a problem for the life-cycle model. It appears that the elderly do not dissave as much as the model predicts. In other words, the elderly do not run down their wealth as quickly as one would expect if they were trying to smooth their consumption over their remaining years of life.

There are two chief explanations for why the elderly do not dissave to the extent that the model predicts. Each suggests a direction for further research on consumption.

The first explanation is that the elderly are concerned about unpredictable expenses. Additional saving that arises from uncertainty is called **precautionary saving**. One reason for precautionary saving by the elderly is the possibility of living longer than expected and thus having to provide for a longer than average span of retirement. Another reason is the possibility of illness and large medical bills. The elderly may respond to this uncertainty by saving more in order to be better prepared for these contingencies.

The precautionary-saving explanation is not completely persuasive, because the elderly can largely insure against these risks. To protect against uncertainty regarding life span, they can buy *annuities* from insurance companies. For a fixed fee, annuities offer a stream of income that lasts as long as the recipient lives. Uncertainty about medical expenses should be largely eliminated by Medicare, the government's health insurance plan for the elderly, and by private insurance plans.

The second explanation for the failure of the elderly to dissave is that they may want to leave bequests to their children. Economists have proposed various theories of the parent-child relationship and the bequest motive. In Chapter 15 we discussed some of these theories and their implications for consumption and fiscal policy.

Overall, research on the elderly suggests that the simplest life-cycle model cannot fully explain consumer behavior. There is no doubt that providing for retirement is an important motive for saving, but other motives, such as precautionary saving and bequests, appear important as well.³

16-4 Milton Friedman and the Permanent-Income Hypothesis

In a book published in 1957, Milton Friedman proposed the **permanent-income hypothesis** to explain consumer behavior. Friedman's permanent-income hypothesis complements Modigliani's life-cycle hypothesis: both use

³ To read more about the consumption and saving of the elderly, see Albert Ando and Arthur Kennickell, "How Much (or Little) Life Cycle Saving Is There in Micro Data?" in Rudiger Dornbusch, Stanley Fischer, and John Bossons, eds., *Macroeconomics and Finance: Essays in Honor of Franco Modigliani* (Cambridge, Mass.: MIT Press, 1986); and Michael Hurd, "Research on the Elderly: Economic Status, Retirement, and Consumption and Saving," *Journal of Economic Literature* 28 (June 1990): 565–589.

Irving Fisher's theory of the consumer to argue that consumption should not depend on current income alone. But unlike the life-cycle hypothesis, which emphasizes that income follows a regular pattern over a person's lifetime, the permanent-income hypothesis emphasizes that people experience random and temporary changes in their incomes from year to year.⁴

The Hypothesis

Friedman suggested that we view current income Y as the sum of two components, **permanent income** Y^P and **transitory income** Y^T . That is,

$$Y = Y^P + Y^T.$$

Permanent income is the part of income that people expect to persist into the future. Transitory income is the part of income that people do not expect to persist. Put differently, permanent income is average income, and transitory income is the random deviation from that average.

To see how we might separate income into these two parts, consider these examples:

- Maria, who has a law degree, earned more this year than John, who is a high-school dropout. Maria's higher income resulted from higher permanent income, because her education will continue to provide her a higher salary.
- Sue, a Florida orange grower, earned less than usual this year because a freeze destroyed her crop. Bill, a California orange grower, earned more than usual because the freeze in Florida drove up the price of oranges. Bill's higher income resulted from higher transitory income, because he is no more likely than Sue to have good weather next year.

These examples show that different forms of income have different degrees of persistence. A good education provides a permanently higher income, whereas good weather provides only transitorily higher income. Although one can imagine intermediate cases, it is useful to keep things simple by supposing that there are only two kinds of income: permanent and transitory.

Friedman reasoned that consumption should depend primarily on permanent income, because consumers use saving and borrowing to smooth consumption in response to transitory changes in income. For example, if a person received a permanent raise of \$10,000 per year, his consumption would rise by about as much. Yet if a person won \$10,000 in a lottery, he would not consume it all in one year. Instead, he would spread the extra consumption over the rest of his life. Assuming an interest rate of zero and a remaining life span of 50 years,

⁴ Milton Friedman, *A Theory of the Consumption Function* (Princeton, N.J.: Princeton University Press, 1957).

consumption would rise by only \$200 per year in response to the \$10,000 prize. Thus, consumers spend their permanent income, but they save rather than spend most of their transitory income.

Friedman concluded that we should view the consumption function as approximately

$$C = \alpha Y^P,$$

where α is a constant that measures the fraction of permanent income consumed. The permanent-income hypothesis, as expressed by this equation, states that consumption is proportional to permanent income.

Implications

The permanent-income hypothesis solves the consumption puzzle by suggesting that the standard Keynesian consumption function uses the wrong variable. According to the permanent-income hypothesis, consumption depends on permanent income; yet many studies of the consumption function try to relate consumption to current income. Friedman argued that this *errors-in-variables problem* explains the seemingly contradictory findings.

Let's see what Friedman's hypothesis implies for the average propensity to consume. Divide both sides of his consumption function by Y to obtain

$$APC = C/Y = \alpha Y^P/Y.$$

According to the permanent-income hypothesis, the average propensity to consume depends on the ratio of permanent income to current income. When current income temporarily rises above permanent income, the average propensity to consume temporarily falls; when current income temporarily falls below permanent income, the average propensity to consume temporarily rises.

Now consider the studies of household data. Friedman reasoned that these data reflect a combination of permanent and transitory income. Households with high permanent income have proportionately higher consumption. If all variation in current income came from the permanent component, the average propensity to consume would be the same in all households. But some of the variation in income comes from the transitory component, and households with high transitory income do not have higher consumption. Therefore, researchers find that high-income households have, on average, lower average propensities to consume.

Similarly, consider the studies of time-series data. Friedman reasoned that year-to-year fluctuations in income are dominated by transitory income. Therefore, years of high income should be years of low average propensities to consume. But over long periods of time—say, from decade to decade—the variation in income comes from the permanent component. Hence, in long time-series, one should observe a constant average propensity to consume, as in fact Kuznets found.

CASE STUDY**The 1964 Tax Cut and the 1968 Tax Surcharge**

The permanent-income hypothesis can help us to interpret how the economy responds to changes in fiscal policy. According to the *IS–LM* model of Chapters 10 and 11, tax cuts stimulate consumption and raise aggregate demand, and tax increases depress consumption and reduce aggregate demand. The permanent-income hypothesis, however, predicts that consumption responds only to changes in permanent income. Therefore, transitory changes in taxes will have only a negligible effect on consumption and aggregate demand. If a change in taxes is to have a large effect on aggregate demand, it must be permanent.

Two changes in fiscal policy—the tax cut of 1964 and the tax surcharge of 1968—illustrate this principle. The tax cut of 1964 was popular. It was announced to be a major and permanent reduction in tax rates. As we discussed in Chapter 10, this policy change had the intended effect of stimulating the economy.

The tax surcharge of 1968 arose in a very different political climate. It became law because the economic advisers of President Lyndon Johnson believed that the increase in government spending from the Vietnam War had excessively stimulated aggregate demand. To offset this effect, they recommended a tax increase. But Johnson, aware that the war was already unpopular, feared the political repercussions of higher taxes. He finally agreed to a temporary tax surcharge—in essence, a one-year increase in taxes. The tax surcharge did not have the desired effect of reducing aggregate demand. Unemployment continued to fall, and inflation continued to rise.

The lesson to be learned from these episodes is that a full analysis of tax policy must go beyond the simple Keynesian consumption function; it must take into account the distinction between permanent and transitory income. If consumers expect a tax change to be temporary, it will have a smaller impact on consumption and aggregate demand.

16-5 Robert Hall and the Random-Walk Hypothesis

The permanent-income hypothesis is based on Fisher’s model of intertemporal choice. It builds on the idea that forward-looking consumers base their consumption decisions not only on their current income but also on the income they expect to receive in the future. Thus, the permanent-income hypothesis highlights that consumption depends on people’s expectations.

Recent research on consumption has combined this view of the consumer with the assumption of rational expectations. The rational-expectations assumption states that people use all available information to make optimal forecasts

about the future. As we saw in Chapter 13, this assumption can have profound implications for the costs of stopping inflation. It can also have profound implications for the study of consumer behavior.

The Hypothesis

The economist Robert Hall was the first to derive the implications of rational expectations for consumption. He showed that if the permanent-income hypothesis is correct, and if consumers have rational expectations, then changes in consumption over time should be unpredictable. When changes in a variable are unpredictable, the variable is said to follow a **random walk**. According to Hall, the combination of the permanent-income hypothesis and rational expectations implies that consumption follows a random walk.

Hall reasoned as follows. According to the permanent-income hypothesis, consumers face fluctuating income and try their best to smooth their consumption over time. At any moment, consumers choose consumption based on their current expectations of their lifetime incomes. Over time, they change their consumption because they receive news that causes them to revise their expectations. For example, a person getting an unexpected promotion increases consumption, whereas a person getting an unexpected demotion decreases consumption. In other words, changes in consumption reflect “surprises” about lifetime income. If consumers are optimally using all available information, then they should be surprised only by events that were entirely unpredictable. Therefore, changes in their consumption should be unpredictable as well.⁵

Implications

The rational-expectations approach to consumption has implications not only for forecasting but also for the analysis of economic policies. *If consumers obey the permanent-income hypothesis and have rational expectations, then only unexpected policy changes influence consumption. These policy changes take effect when they change expectations.* For example, suppose that today Congress passes a tax increase to be effective next year. In this case, consumers receive the news about their lifetime incomes when Congress passes the law (or even earlier if the law’s passage was predictable). The arrival of this news causes consumers to revise their expectations and reduce their consumption. The following year, when the tax hike goes into effect, consumption is unchanged because no news has arrived.

Hence, if consumers have rational expectations, policymakers influence the economy not only through their actions but also through the public’s expectation of their actions. Expectations, however, cannot be observed directly. Therefore, it is often hard to know how and when changes in fiscal policy alter aggregate demand.

⁵ Robert E. Hall, “Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy* 86 (April 1978): 971–987.

CASE STUDY

Do Predictable Changes in Income Lead to Predictable Changes in Consumption?

Of the many facts about consumer behavior, one is impossible to dispute: income and consumption fluctuate together over the business cycle. When the economy goes into a recession, both income and consumption fall, and when the economy booms, both income and consumption rise rapidly.

By itself, this fact doesn't say much about the rational-expectations version of the permanent-income hypothesis. Most short-run fluctuations are unpredictable. Thus, when the economy goes into a recession, the typical consumer is receiving bad news about his lifetime income, so consumption naturally falls. And when the economy booms, the typical consumer is receiving good news, so consumption rises. This behavior does not necessarily violate the random-walk theory that changes in consumption are impossible to forecast.

Yet suppose we could identify some *predictable* changes in income. According to the random-walk theory, these changes in income should not cause consumers to revise their spending plans. If consumers expected income to rise or fall, they should have adjusted their consumption already in response to that information. Thus, predictable changes in income should not lead to predictable changes in consumption.

Data on consumption and income, however, appear not to satisfy this implication of the random-walk theory. When income is expected to fall by \$1, consumption will on average fall at the same time by about \$0.50. In other words, predictable changes in income lead to predictable changes in consumption that are roughly half as large.

Why is this so? One possible explanation of this behavior is that some consumers may fail to have rational expectations. Instead, they may base their expectations of future income excessively on current income. Thus, when income rises or falls (even predictably), they act as if they received news about their lifetime resources and change their consumption accordingly. Another possible explanation is that some consumers are borrowing-constrained and, therefore, base their consumption on current income alone. Regardless of which explanation is correct, Keynes's original consumption function starts to look more attractive. That is, current income has a larger role in determining consumer spending than the random-walk hypothesis suggests.⁶

⁶ John Y. Campbell and N. Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence," *NBER Macroeconomics Annual* (1989): 185–216; Jonathan Parker, "The Response of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review* 89 (September 1999): 959–973; and Nicholas S. Souleles, "The Response of Household Consumption to Income Tax Refunds," *American Economic Review* 89 (September 1999): 947–958.

16.6 David Laibson and the Pull of Instant Gratification

Keynes called the consumption function a “fundamental psychological law.” Yet, as we have seen, psychology has played little role in the subsequent study of consumption. Most economists assume that consumers are rational maximizers of utility who are always evaluating their opportunities and plans in order to obtain the highest lifetime satisfaction. This model of human behavior was the basis for all the work on consumption theory from Irving Fisher to Robert Hall.

More recently, economists have started to return to psychology. They have suggested that consumption decisions are not made by the ultrarational *homo economicus* but by real human beings whose behavior can be far from rational. The most prominent economist infusing psychology into the study of consumption is Harvard professor David Laibson.

Laibson notes that many consumers judge themselves to be imperfect decisionmakers. In one survey of the American public, 76 percent said they were not saving enough for retirement. In another survey of the baby-boom generation, respondents were asked the percentage of income that they save and the percentage that they thought they should save. The saving shortfall averaged 11 percentage points.

According to Laibson, the insufficiency of saving is related to another phenomenon: the pull of instant gratification. Consider the following two questions:

Question 1: Would you prefer (A) a candy today or (B) two candies tomorrow.

Question 2: Would you prefer (A) a candy in 100 days or (B) two candies in 101 days.

Many people confronted with such choices will answer A to the first question and B to the second. In a sense, they are more patient in the long run than they are in the short run.

This raises the possibility that consumers’ preferences may be *time-inconsistent*: they may alter their decisions simply because time passes. A person confronting question 2 may choose B and wait the extra day for the extra candy. But after 100 days pass, he then confronts question 1. The pull of instant gratification may induce him to change his mind.

We see this kind of behavior in many situations in life. A person on a diet may have a second helping at dinner, while promising himself that he will eat less tomorrow. A person may smoke one more cigarette, while promising himself that this is the last one. And a consumer may splurge at the shopping center, while promising himself that tomorrow he will cut back his spending and start saving more for retirement. But when tomorrow arrives, the promises are in the past, and a new self takes control of the decisionmaking, with its own desire for instant gratification.

These observations raise as many questions as they answer. Will the renewed focus on psychology among economists offer a better understanding of

consumer behavior? Will it offer new prescriptions regarding, for instance, tax policy toward saving? It is too early to say, but without doubt, these questions are on the forefront of the research agenda.⁷

16.7 Conclusion

In the work of six prominent economists, we have seen a progression of views on consumer behavior. Keynes proposed that consumption depends largely on current income. Since then, economists have argued that consumers understand that they face an intertemporal decision. Consumers look ahead to their future resources and needs, implying a more complex consumption function than the one that Keynes proposed. Keynes suggested a consumption function of the form

$$\text{Consumption} = f(\text{Current Income}).$$

Recent work suggests instead that

$$\begin{aligned} \text{Consumption} \\ = f(\text{Current Income, Wealth, Expected Future Income, Interest Rates}). \end{aligned}$$

In other words, current income is only one determinant of aggregate consumption.

Economists continue to debate the importance of these determinants of consumption. There remains disagreement about, for example, the influence of interest rates on consumer spending, the prevalence of borrowing constraints, and the importance of psychological effects. Economists sometimes disagree about economic policy because they assume different consumption functions. For instance, as we saw in the previous chapter, the debate over the effects of government debt is in part a debate over the determinants of consumer spending. The key role of consumption in policy evaluation is sure to maintain economists' interest in studying consumer behavior for many years to come.

Summary

1. Keynes conjectured that the marginal propensity to consume is between zero and one, that the average propensity to consume falls as income rises, and that current income is the primary determinant of consumption. Studies of household data and short time-series confirmed Keynes's conjectures. Yet studies of long time-series found no tendency for the average propensity to consume to fall as income rises over time.
2. Recent work on consumption builds on Irving Fisher's model of the consumer. In this model, the consumer faces an intertemporal budget constraint

⁷ For more on this topic, see David Laibson, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 62 (May 1997): 443–477; and George-Marios Angeletos, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evidence," *Journal of Economic Perspectives*, 15(3) (Summer 2001): 47–68.

and chooses consumption for the present and the future to achieve the highest level of lifetime satisfaction. As long as the consumer can save and borrow, consumption depends on the consumer's lifetime resources.

3. Modigliani's life-cycle hypothesis emphasizes that income varies somewhat predictably over a person's life and that consumers use saving and borrowing to smooth their consumption over their lifetimes. According to this hypothesis, consumption depends on both income and wealth.
4. Friedman's permanent-income hypothesis emphasizes that individuals experience both permanent and transitory fluctuations in their income. Because consumers can save and borrow, and because they want to smooth their consumption, consumption does not respond much to transitory income. Consumption depends primarily on permanent income.
5. Hall's random-walk hypothesis combines the permanent-income hypothesis with the assumption that consumers have rational expectations about future income. It implies that changes in consumption are unpredictable, because consumers change their consumption only when they receive news about their lifetime resources.
6. Laibson has suggested that psychological effects are important for understanding consumer behavior. In particular, because people have a strong desire for instant gratification, they may exhibit time-inconsistent behavior and may end up saving less than they would like.

KEY CONCEPTS

Marginal propensity to consume	Normal good	Precautionary saving
Average propensity to consume	Income effect	Permanent-income hypothesis
Intertemporal budget constraint	Substitution effect	Permanent income
Discounting	Borrowing constraint	Transitory income
Indifference curves	Life-cycle hypothesis	Random walk
Marginal rate of substitution		

QUESTIONS FOR REVIEW

1. What were Keynes's three conjectures about the consumption function?
2. Describe the evidence that was consistent with Keynes's conjectures and the evidence that was inconsistent with them.
3. How do the life-cycle and permanent-income hypotheses resolve the seemingly contradictory pieces of evidence regarding consumption behavior?
4. Use Fisher's model of consumption to analyze an increase in second-period income. Compare the case in which the consumer faces a binding borrowing constraint and the case in which he does not.
5. Explain why changes in consumption are unpredictable if consumers obey the permanent-income hypothesis and have rational expectations.
6. Give an example in which someone might exhibit time-inconsistent preferences.

PROBLEMS AND APPLICATIONS

- The chapter uses the Fisher model to discuss a change in the interest rate for a consumer who saves some of his first-period income. Suppose, instead, that the consumer is a borrower. How does that alter the analysis? Discuss the income and substitution effects on consumption in both periods.
- Jack and Jill both obey the two-period Fisher model of consumption. Jack earns \$100 in the first period and \$100 in the second period. Jill earns nothing in the first period and \$210 in the second period. Both of them can borrow or lend at the interest rate r .
 - You observe both Jack and Jill consuming \$100 in the first period and \$100 in the second period. What is the interest rate r ?
 - Suppose the interest rate increases. What will happen to Jack's consumption in the first period? Is Jack better off or worse off than before the interest rate rise?
 - What will happen to Jill's consumption in the first period when the interest rate increases? Is Jill better off or worse off than before the interest-rate increase?
- The chapter analyzes Fisher's model for the case in which the consumer can save or borrow at an interest rate of r and for the case in which the consumer can save at this rate but cannot borrow at all. Consider now the intermediate case in which the consumer can save at rate r_s and borrow at rate r_b , where $r_s < r_b$.
 - What is the consumer's budget constraint in the case in which he consumes less than his income in period one?
 - What is the consumer's budget constraint in the case in which he consumes more than his income in period one?
 - Graph the two budget constraints and shade the area that represents the combination of first-period and second-period consumption the consumer can choose.
 - Now add to your graph the consumer's indifference curves. Show three possible outcomes: one in which the consumer saves, one in which he borrows, and one in which he neither saves nor borrows.
- What determines first-period consumption in each of the three cases?
- Explain whether borrowing constraints increase or decrease the potency of fiscal policy to influence aggregate demand in each of the following two cases:
 - A temporary tax cut.
 - An announced future tax cut.
- In the discussion of the life-cycle hypothesis in the text, income is assumed to be constant during the period before retirement. For most people, however, income grows over their lifetimes. How does this growth in income influence the lifetime pattern of consumption and wealth accumulation shown in Figure 16-12 under the following conditions?
 - Consumers can borrow, so their wealth can be negative.
 - Consumers face borrowing constraints that prevent their wealth from falling below zero.

Do you consider case (a) or case (b) to be more realistic? Why?
- Demographers predict that the fraction of the population that is elderly will increase over the next 20 years. What does the life-cycle model predict for the influence of this demographic change on the national saving rate?
- One study found that the elderly who do not have children dissave at about the same rate as the elderly who do have children. What might this finding imply about the reason the elderly do not dissave as much as the life-cycle model predicts?
- Consider two savings accounts that pay the same interest rate. One account lets you take your money out on demand. The second requires that you give 30-day advance notification before withdrawals. Which account would you prefer? Why? Can you imagine a person who might make the opposite choice? What do these choices say about the theory of the consumption function?